

CSAE WPS/2007-04

Chronic Poverty and All That: The Measurement of Poverty over Time

Cesar Calvo*
Stefan Dercon**

*University of Oxford and Universidad de Piura

**University of Oxford

Revised March 2007

Abstract

We explore how to measure poverty over time, by focusing on trajectories of poverty rather than poverty at a particular point in time. We consider welfare outcomes over a period in time, consisting of a number of spells. We offer a characterization of desirable properties for measuring poverty across these spells, as well as an explicit discussion of three issues. First, should there be scope for compensation so that a poor spell can be compensated for by a non-poor spell? Second, is there scope for discounting or should all spells be equally valued? Third, does the actual sequence of poor spells matter, for example whether they are consecutive or not? We offer a number of measures that implicitly offer different answers to these questions, in a world of certainty. Finally, we also offer an extension towards a forward-looking measure of vulnerability, defined as the threat of poverty over time, that incorporates risk. An application to data from Ethiopia shows that especially the assumption of compensation results in different inference on poverty.

Contacts

:

ccalvo@udep.edu.pe and stefan.dercon@economics.ox.ac.uk

Keywords: poverty, chronic poverty, poverty dynamics, Ethiopia

JEL Codes: D81, I31, I32

1. Introduction

A vast literature has developed on the measurement of poverty. Poverty is considered a state of deprivation, with a living standard below some minimal level. Much debate has focused on ways to approach the underlying standard of living. For example, in recent years much attention has been given to finding appropriate ways to address the multidimensionality in assessing living standards and poverty (Tsui 2002, Bourguignon and Chakravarty 2003). In this paper we focus on another issue often ignored in the standard poverty literature, that the standard of living is not a static, timeless state, but a state that evolves over time. The standard of living follows a trajectory, a path with a history and a future. As a consequence, to assess poverty over time for a particular individual or society, we could explore how we should assess different trajectories of the standard of living, rather than just focusing on the standard of living and poverty in each period, as if neither past nor future poverty experiences had any bearing on the meaning of present hardship. In this paper, we provide some tentative steps to address this issue.

Both theory and empirical evidence provide reasons why careful attention to time paths may be important. If we are interested in wellbeing over a long time span, information about present outcomes can only be sufficient in a very stable world, where individuals need not exert any effort to ensure that their outcomes remain invariant. It is hard to think of such scenario. In practice, first, a myriad of reasons for fluctuations exist, and smoothing efforts are often impossible, e.g. in the case of health, which cannot be transferred from the present to the future, nor vice versa. While some storing technology may be available for other wellbeing dimensions, the individual may still find it hard to fully smooth away all variations, since such technology will rarely be perfect. For instance, in the case of consumption, credit market failures disallow some people to resort to high future consumption flows in the face of current hardship. Secondly, in a world with uncertainty, random shocks may push outcomes above or below the expected time-invariant target. If insurance mechanisms are imperfect, then the individual will be exposed to the consequences of shocks she failed to foresee.

In this paper, fluctuations are interesting in their own right. However, this does not mean that their long-term effects on living conditions are overlooked. Surely enough, fluctuations may turn into serious persistence: a temporary shortfall may translate in a long period of low wellbeing, with slow and uneven recovery, if at all. Also, in their quest for stability, households may react to the threat of fluctuations by resorting to smoothing efforts with some cost in terms of long-run growth. For instance, a street vendor may prefer not to commit to items exhibiting great seasonality, even if they are very profitable.

The issues arising as soon as we pay attention to time trajectories are thus manifold. Policy implications also promptly crop up. For instance, this concern can be directly linked to policy discussion related to concepts of ‘chronic’ poverty: we should be concerned with poverty that does not easily resolve itself, that has a persistence attached to it. Obviously, this is a statement about a future state, but not just about one future period, but related to a permanent escape or the lack of escape from poverty, persisting in different periods. In order to assess different paths over time, means of ordering and/or valuing these trajectories are required.

This paper therefore explores issues related to the assessment of poverty over a lengthy *period* of time for an individual. By ‘lengthy’ we mean that this period can be decomposed into *spells*. In each spell, we observe the level of the standard of living, which for simplicity we will call consumption. Each spell is long enough for consumption flows to be observed and measured. For instance, we may think of a five-year *period*, with consumption data for each single year. Let us use ‘spell’ to refer to the time-units (indexed by t) where consumption flows c_t are measured (in the example, one year), and ‘period’ to refer to that ‘lengthy’ stretch which we are interested in (i.e. all five years together).

While for the sake of concreteness, in this note we prefer to speak of consumption, the discussion equally applies to any other dimension of well-being, such as nutritional status. Define poverty in a T -spell period as

$$P_T(y_1, y_2, \dots, y_T),$$

where y_t stands for consumption at spell t . Let z be the poverty line. We assume this line to be time-invariant for simplicity. Alternatively, if poverty lines did vary over time, our analysis would still hold only if outcomes in every spell to be normalised with respect to their spell-specific poverty lines. Put it differently, in our setting, consumption changes over time must reflect variations in the ability of the individual to reach decent living standards, above the minimum acceptable norm.

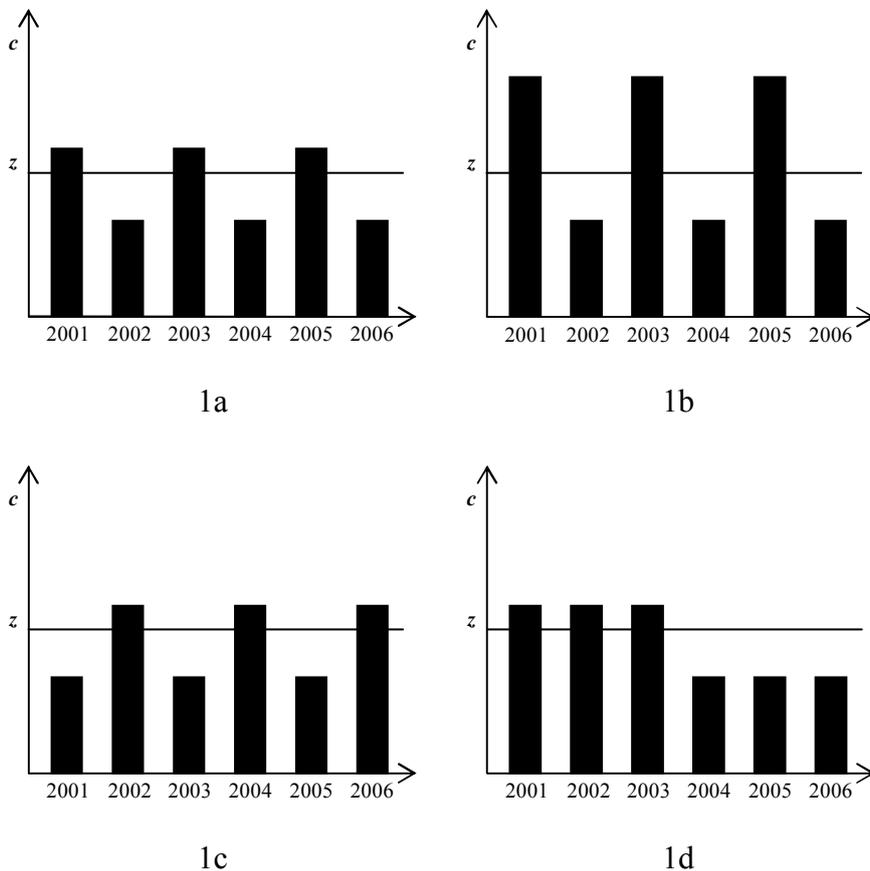
It would be wrong to suggest that the concerns addressed have no precedents. A vast empirical literature has developed that assesses the ‘dynamics’ of poverty, by following the poverty status over time of particular individuals or groups. For example, Baulch and Hoddinott (2000) summarize a number of studies, using panel data, by counting ‘poverty spells’, whereby they mean how often people are observed to be poor in a particular period, and also using simple concepts of poverty mobility, based on poverty transition matrices, identifying who moved in and out of poverty, and who stayed poor. The best-known summarizing measure of poverty assessed over time is Ravallion’s ‘chronic poverty’ measure. This measure assesses chronic poverty as the level of poverty obtained based on a Foster-Greer-Thorbecke measure, using the average level of consumption over the entire period as the underlying standard of living measure (Jalan and Ravallion, 2000).

In this paper, we will argue that these approaches are particular, certainly suggestive, but still arbitrary choices among many different others that could be made to make sense of poverty over a particular period.¹ We will present a number of measures and document some of the specific underlying normative choices based on specific alternative axioms. Our approach may be best motivated by considering a few imaginary scenarios. First, let Figure 1a act as a benchmark description of consumption flows of a given individual. There, each poverty-free spell is succeeded by hardship, which in turn lasts for only one spell and is followed by a fresh episode of sufficient consumption. How should this scenario compare to those in the other three charts? In Figure 1b, the same pattern exists, except consumption is higher in non-poor spells, whereas poor episodes remain just as bad. Should we say that period-

¹ Some of the concerns explicitly considered in this paper related to compensation over time and discounting are also discussed in a very different context, related to adjusting poverty measures to handle differential mortality across a population in Kanbur and Mukherjee (2006).

long poverty has lessened? This raises the issue of *compensation* of poverty spells by non-poverty spells, and the first issue tackled below. As we will show, different plausible measures of poverty of time take a different stance on this issue. In static poverty measurement, across individuals, the issue barely arises by using the focus axiom: the non-poor's outcomes are considered as if they just have reached the poverty line. When considering the poverty over time of a specific individual, this is not self-evidently resolved, as some may argue that hardship at some point in life may be acceptable if it is followed by much better outcomes in other periods. In our measures, we will show that how such judgements can be incorporated.

Figure 1
Illustrative examples



Next, compare 1a with 1c. As seen from 2006, the salient difference lies now in the fact that poverty episodes were suffered further back in the past. The alternation pattern is otherwise still in place. The question is then whether the assessment of period-long poverty must pay the same attention and attach the same weight to all

isolated poverty spells, regardless of how far in the past each occurred. This may be the case if the affliction of human deprivation is seen as an irremediable loss, but on the other hand, its burden can also be imagined to die out as time passes. This is a second issue explicitly discussed: is there any case for using ‘discount rates’, judgements on the relative importance of the present relative to the future or past?

The same question arises as we lastly take Figure 1d. Keeping Figure 1a as the benchmark, 2002 and 2005 seem to swap consumption levels. However, a new issue comes forth, since poverty spells are now contiguous, and the individual faces a prolonged episode of poverty (2004-2006). Should the distress of hardship compound over time, such that three-spell episode of poverty should cause greater harm than three isolated poverty spells? This is the third question to tackle as we turn to our intent to propose poverty measures over a lengthy period.

While these stylized examples show some of the choices involved, trajectories observed in actual data look more messy. For example, take four trajectories found in the Ethiopian rural household panel data survey, with six observations in the period 1994 to 2004. While these consumption levels may well be measured with error, the patterns are not simple, and general judgements about how to order these in terms of poverty over time are not self-evident. For example, the household of Abebe (Figure 2a) appears to have been going downhill in the last four years of the data, but has only one spell in poverty, while Alemu (Figure 2b) has four poverty spells, but by the end of the period has two years above the poverty line. Tigist and Asfaw’s families (Figures 2c and 2d) both have spells below the poverty line, but at different times in the sequence

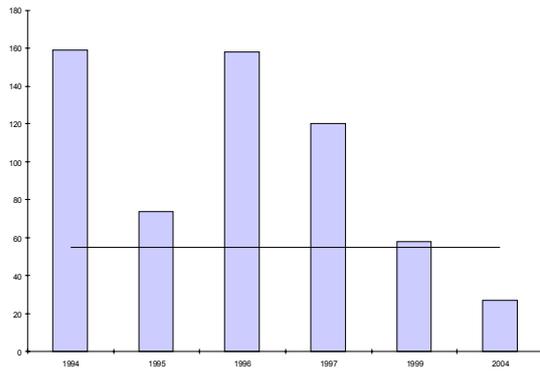


Figure 2a Abebe

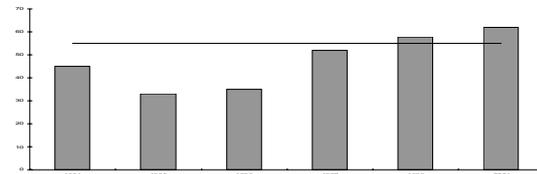


Figure 2b Alemu

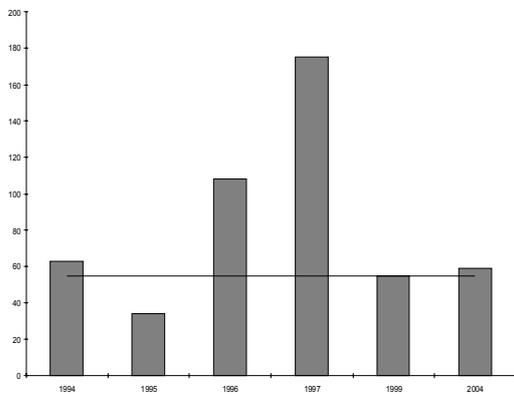


Figure 2c Tigist

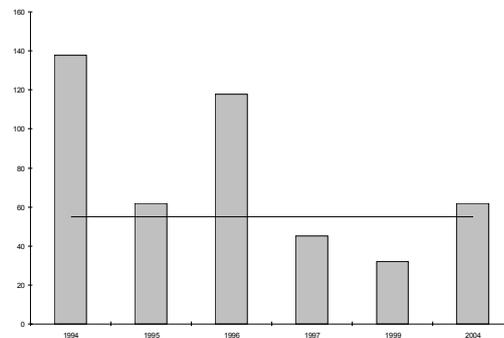


Figure 2d Asfaw

It is clear that many judgements will be required to summarize such trajectories of the standard of living in one single index of intertemporal poverty. This paper aims to present a number of possible indices, even though its main aim is to make some of these normative judgements explicit.

Foster (2007) in this collection has a related objective, but aims to explicitly construct a class of measures of ‘chronic poverty’. Below we will highlight the similarities with at least one of our own measures, but there is one crucial difference worth commenting on now. His measure starts from the identification within the data of who is chronically poor, and then proceed in ways not dissimilar to ours. In his paper, a ‘chronic’ poor person is someone who experiences at least a specific percentage of poverty over time. His measure of chronic poverty then values the depth and severity of poverty for such persons, excluding the non-chronically poor. While internally

fully consistent and sensible, and a chronic poverty equivalent of the Foster-Greer-Thorbecke measure, one key requirement is a judgement of a cut-off for classifying someone as ‘chronic’ poor, irrespective of how far below the poverty line this person is. By introducing a further threshold beyond the poverty line, the result is that people with just over the required number of spells for chronic poverty, but with all spells just below the poverty line would be considered chronically poor, while someone with marginally fewer but more serious spells is counted as transitory poor. Our approach does not resolve this issue at all – it just ducks it – by considering measures of ‘a poor life’, or more precisely the extent, depth and severity of ‘poverty over a period of time’, using a means of weighing all poverty spells in one aggregate across time, irrespective of the frequency of spells.

In the next section, we offer the basic setup, discussing the key decision needed regarding applying a focus on poverty, transformation of outcomes and aggregation over time. In section 3, we present a number of core axioms that may guide these choices, and the resulting choice of measures. In section 4, a set of measures are presented, ordered by the particular sequence in terms of applying focus, transformation and aggregation. They can be shown to satisfy (or not) some of the suggested axioms for intertemporal poverty assessment. The rest of the paper will offer extensions. In section 5, a discussion is introduced on the role of time preference, while in section 6, the idea of a sensitivity to prolonged poverty is introduced. In section 7, we reintroduce risk and derive a forward-looking measure of the threat of long-term poverty, building on our previous work on vulnerability. Finally, in section 8, we offer some examples of how some of the measures may be applied using data from Ethiopia.

2. Basic Setup

Unless otherwise stated, and for most of the paper, we will imagine the world to be uncertainty-free. All consumption levels are perfectly known, regardless of the point we take in time. For instance, as seen from the final spell, a backward-looking assessment of poverty throughout the period has the benefit of hindsight, and no uncertainty clouds the view of past consumption levels. Our assumption intends to put

forward-looking assessments on a similar standing, by granting the individual the gift of perfect foresight. To see what this implies, imagine periods are seen (*ex-post*) from their final spells and ranked according to some intertemporal poverty measure, and also, that some ranking reshuffling occurs if the standing point is brought forward (*ex-ante*) to their first spells. In our world, uncertainty cannot act as an explanation for such reshuffling, at least for now. In the final section of the paper, we will suggest an extension in which this perfect foresight is dropped and uncertainty is reintroduced.

By assuming away uncertainty, we can focus both *ex-post* and *ex-ante* analyses on our central question, which is to identify a metric for *how much suffering or deprivation was or will be endured over a particular period*. This concern must be distinguished from the *current experience* of suffering which may be caused by a grim *future* (a sense of hardship to come), or by unhappy memories of *past* deprivation. For instance, if we speak about poverty between 2007 and 2015, we will enquire how much poverty will be ‘accumulated’ by the end of 2015 (and not how much future hardship impinges *ex-ante* on wellbeing in 2007). In this note, we think of period-poverty as the cumulative result of spell-specific poverty episodes.

In the vein of the distinction between ‘identification’ and ‘aggregation’ in the measurement of aggregate poverty (as in Sen), let us propose the following three stages for our analysis:

a) *Focus*

It is well-known that all measures of aggregate poverty (e.g. Foster, Greer and Thorbecke, 1984) build on some form of focus axiom, whereby outcomes above z are censored down to the poverty line itself, since the poverty of the poor is not meant to be alleviated by the richness of others. For instance, a society will not be said to be less poor simply because the rich become richer, with no change in consumption levels among the poor. Thus, it is this focus condition what ‘identifies’ the relevant outcomes. Let this stage be related to a function $f(u)$, such that $f(u) \equiv \text{Min}[u, z_u]$, where z_u is the relevant poverty threshold, e.g. $z_u = z$ if $y = u$.

b) Transformation

To motivate this stage, recall the well-known Pigou-Dalton condition, whereby aggregate poverty rises if consumption is transferred from the very poor to the not-so-poor. In our case, we may require period poverty P_T to rise as a consequence of a transfer from a poor spell to a not-so-poor spell, in the presumption that the drop in the former will outweigh the gain in the latter. This however is a presumption that cannot be taken for granted in our case, since the locations in time of these two spells may matter and have not been determined yet. For instance, the poorer spell may have occurred such a long time ago, that its loss in consumption may be meaningless. The ‘aggregation’ step turns to such issues shortly.

Nonetheless, we can still say that *for equally-valued spells* (in the way that all individuals are equally valued by the Pigou-Dalton condition), a transfer for the benefit of a not-poor spell should result in greater intertemporal poverty P_T . In practical terms, this implies that outcomes y_t must at some point be transformed by a suitable strictly convex function, *either before or after some correction for the value of their time location has been made.*

Let function $g(u)$, with $g'(u) < 0$ and $g''(u) > 0$ account for the possibility of this strictly convex transformation.

c) Aggregation

Thirdly, spell-specific inputs must combine into one single measure of total, period-long poverty P_T . To be clear, we deal here with an aggregation over time-spells, and not over individuals (as in the usual poverty measures). Hence, aggregation methods may well differ from the standard procedure. For instance, they will need to account for weight differences across time spells, e.g. if we were to decide that spells further back in the past should be paid less attention than more recent ones.

Let function $A_T(u_1, u_2, \dots, u_T)$ perform this aggregation. To keep the convex, Pigou-Dalton-like transformation as a separate issue, let A_T be *linear in each of its arguments*. This restriction has no major drawback – except for the transfer argument described above, there is no obvious reason why changes in any spell should be allowed to have any bearing on the effect on P_T of further changes in that same spell.

How these three stages come together is a question with no unique answer. Will we first apply focus, then transform and then aggregate over the entire period, or will we change the order of these actions? As will be shown below, this sequence matters. But which order we choose will depend on our view on the set of desirable properties of a period-long poverty measure. To develop this further, we will in section 4 give examples of the possible permutations related to focus, transformation and aggregation. In the next section, we will first discuss some possible desiderata.

3. Formalizing the axioms

In this section, we offer a few possible axioms that can guide us in choosing particular measures of individual, period-long poverty. The set of these axioms is not exhaustive, in the sense that no combination of them determines uniquely a particular family of measures. These desiderata will nevertheless offer routes to decide among different permutations of focus, transformation and aggregation.

The first two axioms are quite general and hardly debatable.

Monotonicity in outcomes

Since consumption rises can under no circumstances cause a rise in poverty, we impose

$$\text{For } d > 0, P_T(y_1, y_2, \dots, y_t + d, \dots, y_T) \leq P_T(y_1, y_2, \dots, y_t, \dots, y_T) \quad (1)$$

A narrower definition specification is only possible if we decide when the focus stage will enter. For instance, if the focus function f is allowed to come first, then we could go further and require

$$\text{For } d>0 \text{ and } y_t < z, P_T(y_1, y_2, \dots, y_t + d, \dots, y_T) < P_T(y_1, y_2, \dots, y_t, \dots, y_T) \quad (1')$$

In words, this alternative version imposes that a consumption rise during a poverty spell will reduce overall poverty. In (1), since focus has not yet been enforced, $y_t < z$ is not enough to take for certain that the reduction in P_T will occur, and we can only rule out a rise in poverty.

Increasing cost of hardship

This axiom echoes the Pigou-Dalton condition, whose exact translation to our setting would impose that a consumption transfer from a very-poor spell to a not-so-poor spell should raise overall poverty. However, we cannot readily resort to this formulation here, since the time location of these spells must also be specified, unless we assume that regardless of these locations, all spells are equally valued. While the Pigou-Dalton assumption that all individuals (with equal consumption) receive equal attention faces no major objection, here we must allow for the case where some time spells receive greater weight than others.

An alternative formulation can build on the effect of consumption changes at *one single spell*, and thus steer clear of the risk of committing to valuations of changes in two different spells. The spirit of this condition remains unchanged – *consumption losses hit harder if consumption is already low to begin with*. We may phrase it as the increasing cost of hardship. Formally,

$$\begin{aligned} &\text{For } d>0 \text{ and } y_k < z, P_T(y_1, \dots, y_k, \dots, y_T) - P_T(y_1, \dots, y_k + d, \dots, y_T) \\ &> P_T(y_1, \dots, y_k + d, \dots, y_T) - P_T(y_1, \dots, y_k + 2d, \dots, y_T) \end{aligned} \quad (2)$$

Next, one may invoke an axiom providing the basis for comparison across periods of different lengths. While ‘total’ poverty over a given T -span is necessarily dependent on its length T , one may wish to speak of poverty at an ‘average’ spell, i.e. the spell-

specific poverty level which, if repeated in every single spell of the period, would lead to the observed period-long poverty level. To formalize this, let $P_T(y_1, y_2, \dots, y_T)$ increase proportionally to a k -fold repetition of the period at hand, which we may write as

Full-period repetitions

$$P_{kT}(y'_1, y'_2, \dots, y'_{kT}) = kP_T(y_1, y_2, \dots, y_T),$$

where $y'_{i+k(t-1)} = y_t$ for $t=1, 2, \dots, T$ and $i=1, 2, \dots, k$. (3)

In (3), we imagine that the complete period is lengthened by allowing the first spell repeat k times before the outcome of the (initially) second spell obtains, which then repeats k times before the third outcome occurs, and so forth. Consider the following alternative formulation (3'), where the whole period unfolds and is then followed by an identical sequence of T spells, and then by another, and so forth until it is repeated k times:

$$P_{kT}(y'_1, y'_2, \dots, y'_{kT}) = kP_T(y_1, y_2, \dots, y_T),$$

where $y'_{i+k(i-1)} = y_t$ for $t=1, 2, \dots, T$ and $i=1, 2, \dots, k$ (3')

The difference between (3) and (3') is trivial only if we impose two assumptions which we shall discuss further on, namely that all outcomes are equally valued, regardless of the time when they occur, and also that hardship is assessed in each spell separately, e.g. with no chance for the immediately preceding outcomes to matter. Otherwise, if either of these assumptions fails, then a choice between (3) and (3') is required. Both assumptions are also underlying the two remaining axioms of this section.

Note that this axiom is clearly akin to the population invariance axiom of aggregate poverty measurement. It also plays a similar role here in contributing to a linear specification of the aggregation function A_T . In the spirit of Foster and Shorrocks (1991), we will approach linearity by imposing this full-period repetitions axiom and also a 'sub-period consistency' axiom. The latter is meant to impose Gorman separability and hence needs $P_T(y_1, y_2, \dots, y_T)$ to be a transform of a linear combination

of y_1, y_2, \dots, y_T as long as $T \geq 3$. The axiom on full-period repetitions then generalises this result to $T \geq 1$.

To formalize, and further mirroring the aggregate poverty literature, allow the period to be decomposed into (any) two sub-periods and focus on the reaction of period-long poverty to changes in these sub-periods.

Sub-period consistency

$$\begin{aligned}
 P_T(y'_1, y'_2, \dots, y'_K, y'_{K+1}, \dots, y'_T) &> P_T(y_1, y_2, \dots, y_K, y_{K+1}, \dots, y_T) \\
 \text{if } P_K(y'_1, y'_2, \dots, y'_K) &> P_K(y_1, y_2, \dots, y_K) \\
 \text{and } P_{T-K}(y'_{K+1}, y'_{K+2}, \dots, y'_T) &= P_{T-K}(y_{K+1}, y_{K+2}, \dots, y_T).
 \end{aligned} \tag{4}$$

If some sub-period exhibits a rise in poverty (while poverty remains unaltered in all other sub-periods), then P_T must also rise for the entire period. This sensitivity is what we mean by ‘consistency’. Its interpretation may gain from noting that it restricts the ability of some spells (those from $K+1$ to T) to impinge on the effect of other spells (from 1 to K) on P_T . We may see the seed of a linear specification here, which however needs a stronger axiom to be fully imposed. Such axiom can be phrased as ‘sub-period decomposability’, whereby total period-poverty is a weighted sum of both sub-period poverty indices.

Sub-period decomposability

$$P_T(y_1, y_2, \dots, y_K, y_{K+1}, \dots, y_T) = \frac{K}{T} P_K(y_1, y_2, \dots, y_K) + \frac{(T-K)}{T} P_{T-K}(y_{K+1}, y_{K+2}, \dots, y_T) \tag{5}$$

Needless to say, this axiom is reminiscent of sub-group decomposability in the aggregate poverty literature. Again, note both that timing is assumed to have no bearing on the valuations of a given spell, and any information in the *sequence* of poverty spells can be ignored: the valuation of a poverty spell is unrelated to its history, such as whether the person was poor before or not – sequences are quite freely broken into sub-pieces.

Even though this set of axioms is relatively limited, they are enough for a clarifying discussion on some possible measures, linked to particular permutations of the choices related to focus, transformation and aggregation.

4. Choices as a matter of sequencing

We said that alternative P_T specifications follow from alternative sequencing choices for three crucial stages (Focus, Transformation and Aggregation). Even though six orderings thereof are possible (FTA, TFA, FAT, TAF, AFT and ATF), in this section we only consider four of them, prior to giving the general specification of the corresponding period-long poverty measures, as well as a number of specific examples. Four orderings are enough to characterise the existing alternatives, since it can be easily shown that focus and transformation can swap positions with no practical consequence, provided aggregation is not inserted between them. Thus, FTA exhausts all the insights in TFA, and likewise AFT can stand for ATF.

Case 1: Focus-Transformation-Aggregation (FTA)

$$P_T(y_1, y_2, \dots, y_T) \equiv A_T(g(f(y_1), f(y_2), \dots, f(y_T))) \quad (6)$$

In this case, imposing first focus implies that consumption levels are immediately censored. Hence, this specification rules out compensations *across time spells*, in the same spirit of the focus axiom in aggregate poverty measures, which discards compensations across individuals. In our case, the intuition could be phrased as follows: ‘poverty episodes cause shock and distress to such an extent, that they leave an *indelible mark* – no future or past richness episode can make up for them’.

Under FTA, convexity is imposed next, before aggregation. Unsurprisingly, the resulting families of measures are reminiscent of the well-known Foster, Greer and Thorbecke (1984), and Chakravarty (1983) measures of aggregate poverty. To see this, consider in particular the first two of the examples below, where $\tilde{y}_i \equiv \text{Min}[z, y_i]$ and

aggregation allows for some time-adjustment, as by β^t . For now, and for the rest of this section, we may take this factor as given, until we turn to discuss it in section 4.

$$P_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} \left(1 - \frac{\tilde{y}_t}{z}\right)^\alpha, \text{ with } 1 < \alpha \text{ and } \beta > 0. \quad (7)$$

$$P_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} \left[1 - \left(\frac{\tilde{y}_t}{z}\right)^\alpha\right], \text{ with } 0 < \alpha < 1 \text{ and } \beta > 0. \quad (8)$$

Measure (7) is a simple multi-period version of the FGT measure, where aggregation has acted upon time spells, rather than individuals. (8) offers a similar idea for the Chakravarty measure. In terms of formal desiderata, FTA does rather well in capturing at least some of the basic desiderata. *Both Monotonicity and Increasing cost of hardship apply.* Transformation ensures the latter because it applies before aggregation, i.e. before all spell-specific outcomes merge into some form of total consumption, where no distinction between poor and not-so-poor spells would be possible. Likewise, *Sub-period decomposability* is also possible due to the fact that aggregation comes last, so that the linearity of the final specification is not endangered – thus, total-period poverty can be written as a weighted average of sub-period poverties. This equally allows *Full-period repetitions*.

A limiting case of (7) is familiar, imposing $\beta=1$ and $\alpha=0$. It would result in a period-long poverty measure that simply counts the number of spells below the poverty line. But unlike (7), by imposing $\alpha=0$, it would fail both the *Monotonicity* and the *Increasing cost of hardship* axioms. Nonetheless, the simplicity of this specification makes it a useful starting point for summarizing total-period poverty. It has been used among others by Baulch and Hoddinot (2000), when counting poverty spells and its distribution across a population.

Measure (7) is probably the most straightforward and relevant for empirical analysis. It aggregates individual period-by-period poverty spells into one aggregate measure of poverty over a period of time consisting of T spells. It is also close to the ideas behind Foster (2007) in this collection, with one crucial differences: we do not restrict this

measure to be zero for those who experience a frequency of spells below the ‘chronic poverty’ threshold.

Case 2: Focus-Aggregation-Transformation (FAT)

$$P_T(y_1, y_2, \dots, y_T) \equiv g(A_T(f(y_1), f(y_2), \dots, f(y_T))) \quad (9)$$

A different set of families obtains if aggregation occurs before a convex transformation is enforced. Since focus retains the first move, it is still true that poverty episodes remain the crucial concern. We do not take into account any outcomes above the poverty line: there is no weight attached to being better off in good years. For example, Figures 1a and 1b will still be equally valued and period-long poverty will still be the same for both cases. However, in (9), the severity of poverty is paid attention to not in every single spell, but only after all spell-specific outcomes are summarised into one single value. *It is overall severity that matters.*

Take the following few example, which can read as a transformation of some form of ‘present value of censored consumption’:

$$P_T(y_1, y_2, \dots, y_T) = \left(\sum_{t=1}^T \beta^{T-t} \left[1 - \frac{\tilde{y}_t}{z} \right] \right)^\alpha, \text{ with } 1 < \alpha \text{ and } 0 < \beta. \quad (10)$$

Given the sequencing of the three stages, it is clear that *Monotonicity* still holds, unlike *Sub-period decomposability* and *Full-period repetitions*, which must be weakened down to *Sub-period consistency*. Finally, *Increasing cost of hardship* also fails to hold, which may be undesirable on a number of accounts – very bad poverty spells are brushed aside as long as on average, poverty spells are not too severe. This result, which clearly follows from the location of transformation at the final position of the sequence, may explain why no instances of this specification can be found in the literature. Nonetheless, other cases where transformation also comes last do exist in the literature, as we see next.

Case 3: Aggregation-Focus-Transformation (AFT)

$$P_T(y_1, y_2, \dots, y_T) \equiv g(f(A_T(y_1, y_2, \dots, y_T))) \quad (11)$$

Here, transformation remains last, and even more importantly, focus is removed from the first position. Note that this second choice implies that some degree of *compensation does occur across spells*. As opposed to the view underlying FTA and FAT, what matters here is not so much whether the individual faced severe hardship at any particular point in time (regardless of how she performs at other points). The main concern is rather that outcomes realised in the rest of the period may not be high enough to compensate for observed hardship episodes. In other words, *poverty does not imply an irremediable loss*, since the case is also possible, where hardship does occur, but high consumption in other spells do ‘save’ the period. Looking back at our illustrative examples, Figure 1a has more poverty than Figure 1b.

Put it differently, (11) would be consistent with poverty assessed in relation to some form of intertemporal utility-based measure of poverty, whereby, given instantaneous or direct utility in a particular spell, the present value of these utilities is calculated as the sum of discounted direct utility, to which then some benchmark norm is applied. While this is open to argument, it does make somewhat unsatisfactory reading since period-long poverty can be reduced by focusing on spells of already high consumption well above the poverty line, say in the form of temporary opulence and feasts. Nevertheless, some intuitive examples can be shown of measures in this case:

$$P_T(y_1, y_2, \dots, y_T) = \left\{ 1 - \text{Min} \left[1, \frac{\sum_{t=1}^T \beta^{T-t} y_t}{\sum_{t=1}^T \beta^{T-t} z} \right] \right\}^\alpha, \text{ with } 1 < \alpha \text{ and } 0 < \beta. \quad (12)$$

$$P_T(y_1, y_2, \dots, y_T) = 1 - \left\{ \text{Min} \left[1, \frac{\sum_{t=1}^T \beta^{T-t} y_t}{\sum_{t=1}^T \beta^{T-t} z} \right] \right\}^\alpha, \text{ with } 0 < \alpha < 1 \text{ and } 0 < \beta. \quad (13)$$

Since transformation comes last, *Increasing cost of hardship* fails to hold. More strikingly, also *Sub-period consistency* does (which of course rules out *Sub-period decomposability* as well). To see why, take the following example. Imagine outcomes in a four-spell period changes from (8,8,8,40) to (4,4,8,40), with $z=10$. Poverty has

risen in the sub-period comprising the first two spells (while the rest of the period is unaltered), and yet poverty for the entire poverty remains at zero. Again, the reason must be found in the fact that compensations across spells are possible.

This may therefore seem an unappealing measure. However, one of the most commonly used ‘measures of chronic poverty’, based on Jalan and Ravallion (2000), is directly nested in this case, for $\beta=1$. The measure reduces then to

$$P_T(y_1, y_2, \dots, y_T) = \left\{ 1 - \text{Min} \left[1, \frac{\sum_{t=1}^T y_t}{Tz} \right] \right\}^\alpha, \text{ which is an FGT measure of poverty}$$

applied to *mean* consumption in the period. It rests strongly on the case for compensations across periods.

Case 4: Transformation-Aggregation-Focus (TAF)

$$P_T(y_1, y_2, \dots, y_T) \equiv f(A_T(g(y_1), g(y_2), \dots, g(y_T))) \quad (14)$$

Again, removing focus from the first position does matter, since compensations are allowed. For instance in the following examples, the main comparison takes place between the norm and some aggregation of the stream of consumption flows (say, its present value):

$$P_T(y_1, y_2, \dots, y_T) = \text{Max} \left[0, \sum_{t=1}^T \beta^{T-t} \left(1 - \frac{y_t}{z} \right)^\alpha \right], \text{ with } 1 < \alpha \text{ and } 0 < \beta. \quad (15)$$

$$P_T(y_1, y_2, \dots, y_T) = \text{Max} \left[0, \sum_{t=1}^T \beta^{T-t} \left[1 - \left(\frac{y_t}{z} \right)^\alpha \right] \right], \text{ with } 0 < \alpha < 1 \text{ and } 0 < \beta \quad (16)$$

Note that in fact, these specifications may allow $P_T=0$ even if $y_t < z$ for some t – this may well be the case if y_t is sufficiently above z in some other spells.

In terms of our desiderata, *Increasing cost of hardship* applies (since transformation is enforced before aggregation), but again, *Sub-period consistency* is dropped, along with *Sub-period decomposability*. In addition, *Monotonicity* is risked, since cases

where $y_t > z$ will display the troublesome feature of greater *positive* gaps between y_t and z raising both spell-specific and period-long poverty. Unsurprisingly, no instance of this specification exists in the literature.

The result of this discussion is that a number of choices can be made in terms of the sequence of aggregation, transformation and the application of a focus criterion, but only a relatively limited set is consistent with some desiderata. For example, (7) and (8) or (12) and (13), building on the Foster-Greer-Thorbecke and Chakravarty families of measure have been taken as acceptable candidates. A key issue is the extent of compensation between spells that is allowed – a normative choice we can only point to. However, the discussion opens avenues for applications and extensions. In the next few sections, we will address three further issues: first, whether there is any primacy of particular spells in our assessment of period-long poverty. For example, should the last state be given any special weight, as the end-point of our assessment? The second issue is whether there are any normative issues related to the particular sequencing of spells – in particular, should any additional attention be paid to repeated spells and therefore prolonged periods of poverty? Finally, what would happen if we move to forward-looking measures, that take into account that the world is uncertain?

5. Equally-valued spells

In all the examples thus far we have not been explicit about the choice of the parameter β beyond requiring that it is positive. The coefficient β determines the rate of time discounting: the weight we attach to consumption and poverty spells in different periods. Standard economic analysis assesses the value of some future flow of a variable of interest (such as income or consumption) by assuming the rate of time discounting to give nearby flows a higher weight. For instance, in (7), where

$$P_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} \left(1 - \frac{\tilde{y}_t}{z} \right)^\alpha,$$

as seen from the outset at $t=1$, we would require $\beta > 1$, so that outcomes in the distant future receive less attention. However, in our assessment of poverty spells, we argue that in fact, the choice is not as self-

evident. For example taking $\beta=1$, meaning not to exert any discount, or even $\beta<1$, may be a sensible decision, given our purposes.

While time-discounting is made undisputed use of in most intertemporal economic problems, it is not self-evident that it should apply when assessing hardship spells. Severe hardship must cause some irremediable impact on human life, or at least this seems to be the spirit underlying the whole of the literature on poverty. Poverty episodes are spells of misfortune which cannot be compensated for (in the spirit of FTA). Note how close this argument comes to the rationale behind the cases above, where focus is given the first priority, as opposed to those where the focus applies after some aggregation has been performed and outcomes are allowed to compensate for one another across spells, such as in the case of AFT. Even if some compensation were allowed for, it would seem reasonable to require that compensation comes at least at some serious cost. In any case, allowing some compensation is not an argument to dismiss poverty spells, simply because they occur far away in the future. In other words, discounting spells would sit uncomfortable with a concept of period-long poverty.

There is a corollary in the literature on health measures. In the context of the measurement of health, Anand and Hanson (1994) refuse to accept time discounts in the calculation of DALYs: “We can see no justification for an estimation of the time lost to illness or death which depends on when the illness or the calculation occurs. Suppose a person experiences an illness today and another person, identical in all respects, experiences an illness of exactly the same description next year. Discounting amounts to concluding that the quantity of the (same) illness is lower in the latter case. This does not accord with intuition or even with common use of language”. We are inclined to agree with this view. “A principle of universalism would argue strongly for a common intrinsic valuation of human life, regardless of the age at (or the time period in) which it is lived”.

An axiomatic formulation for this stance ($\beta=1$) allows reshuffles across time-positions to occur with no bearing on total, period-long poverty. Timing does not matter. Thus, we could impose

Symmetry over time-positions

$$P_T(y_1, y_2, \dots, y_T) = P_T(y_{\sigma(1)}, y_{\sigma(2)}, \dots, y_{\sigma(T)}), \quad (17)$$

where $\sigma(u)$ is a one-to-one function whose co-domain is identical to its domain $(1, 2, \dots, T)$.

All the measures described before could be trivially adjusted to allow for (17) by setting $\beta=1$.

But other arguments could be made. In evaluating trajectories, one may well be tempted to value more the spells at the end of period rather those at the beginning. Gradually drifting into poverty is then viewed as worse than evolving from spells in poverty out of poverty, even if the number and extent of spells in poverty may be equal in both cases. “All is well that ends well” may be sentiment that could be reflected in our value judgements. An example could be Figure 1d, compared to the reverse of this graph whereby the three ‘non-poor’ spells come at the end: the latter would then be considered better. One way of introducing this in our evaluation of trajectories would be to consider $\beta < 1$: spells later on are given a higher weight. Other choices are also possible: β could become period-dependent and particular periods in the future could be given a much higher weight.²

6. Axioms of a sequence-sensitive specification

An arguably strong assumption is that some form of linearity is always present in our measures of period-long poverty. To be more precise, note that our aggregation function A_T rules out any cross-effect across spells, i.e. $\frac{\partial^2 A_T}{\partial y_s \partial y_t} = 0$. This linearity is at

the basis of the fact that the valuation of a poverty spell is unrelated to its history, such as whether the person was poor before or not. This is ensured by the linearity-related axioms above, but it can also be summarised by an underlying axiom ensuring

² This sentiment is not unknown in the policy discourse, where targets are set: the Millenium Development Goals have a well-defined deadline, 2015, and this deadline is seemingly far more important than, say, outcomes in the preceding years.

Independence of other time-spells

$$\begin{aligned} & P_T(y_1, \dots, y_{K-1}, y'_K, y_{K+1}, \dots, y_T) - P_T(y_1, \dots, y_{K-1}, y_K, y_{K+1}, \dots, y_T) \\ & = P_T(y'_1, \dots, y'_{K-1}, y'_K, y'_{K+1}, \dots, y'_T) - P_T(y'_1, \dots, y'_{K-1}, y_K, y'_{K+1}, \dots, y'_T) \end{aligned} \quad (18)$$

However, the case against such independence exists. Indeed, one may prefer to imagine that prolonged, uninterrupted poverty is less acceptable than a situation of equally-frequent, but intermittent poverty episodes. For instance, within a $T=3$ period, two poverty episodes in a row may be harder to bear than the same two poverty episodes with a recovery-spell in between.

Of course, this is a normative issue. It may also be phrased on the grounds of *technology-related mechanisms*, which we may even provide with the support of some empirical evidence – e.g. body strength is progressively undermined by continuous hardship and makes further poverty harder to bear, or more plainly, low consumption comes hand in hand with asset depletion. However, we prefer to say that prolonged poverty can be particularly bad *per se*.³ The quality of a human life may be eroded more harshly if poverty is sustained for a lengthy string of spells.

In this case, we may define

Prolonged poverty

$$\begin{aligned} & P_T(y_1, y_2, \dots, y_{K-1}, y_K + d, \dots, y_T) - P_T(y_1, y_2, \dots, y_{K-1}, y_K, \dots, y_T) \leq \\ & P_T(y_1, y_2, \dots, y_{K-1} + e, y_K + d, \dots, y_T) - P_T(y_1, y_2, \dots, y_{K-1} + e, y_K, \dots, y_T), \text{ for } d, e \geq 0. \end{aligned} \quad (19)$$

This axiom implies that some form of path dependence exists. A change in any given spell can only be assessed with knowledge of outcomes in previous spells. In particular, greater poverty in a spell implies that a drop in consumption in the

³ Another way of putting this is that we assume here that our underlying standard of living indicator comprehensively incorporates these concerns, so that there is no more information on the spell-specific standard of living required, for example on one's asset position, once the standard of living is known. Our concern with the sequence of poverty spells relates to assessing the sequence of spell-specific standard of living outcomes: repeated spells have an additional welfare cost and there is information in the sequencing of spells.

following spell will hit harder. Our specification in (19), however, can only be taken as a starting point, since it narrows the concept of prolonged poverty down to a dependence only on the immediately preceding spell, whereas one may just as well allow spells further back to matter likewise.

Note that this concern with prolonged poverty is not just one more form of smoothing behaviour. In fact, it may actually run against such behaviour. For instance, in the face of three consecutive spells where the consumption level remains invariant and below the poverty line, $P_T(y_1, y_2, \dots, y_T)$ may drop if the neat, smooth sequence is broken by raising the middle consumption level above the poverty line, at the expense of a decrease in the other two spells. In other words, individual preferences may or may not favour smoothing efforts, and yet sensitivity of $P_T(y_1, y_2, \dots, y_T)$ to prolonged poverty persists all the same. Our measure has a normative role, consisting in no more than reporting the extent of poverty-related suffering over a stretch of time, quite regardless of the features of the objective function of the individual.

For instance, take the following specification:

$$P_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} h(\tilde{y}_t, \tilde{y}_{t-1}), \text{ with } 0 < \beta \quad (20)$$

and where some standard value for \tilde{y}_0 could be added as a convention to prevent $h(\tilde{y}_t, \tilde{y}_{t-1})$ from being undefined for $t=1$. One particular specification for this sequence-sensitive measure could be

$$P_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} \left(\left(1 - \frac{\tilde{y}_t}{z} \right) \left(1 - \frac{\tilde{y}_{t-1}}{z} \right)^\rho \right)^\alpha, \text{ with } 1 < \alpha, 0 < \beta \text{ and } 1 < \rho < 0 \quad (21)$$

The measure in (21) can be seen as one example of a FTA measure: first focus is applied, then a transformation takes place and finally, aggregation over spells. This last stage includes however a new element, as it allows the preceding spell to act as a weight. In particular, note that quite naturally a poverty-free spell ($\tilde{y}_t = z$) does not add to period-long poverty, neither in the same spell nor in the following one. However, if poverty does hit the individual, then the resulting burden increases in the severity of

hardship in the recent past (since $\rho > 0$). And likewise, this new poverty episode impinges on the weight of future deprivation. The restriction $\rho > 1$ simply aims to rule out the case where in the assessment of hardship at time t , the poverty gap at $t-1$ receives more attention than the actual gap at t .

Note that this specification imposes that \tilde{y}_t and \tilde{y}_{t-1} must be seen as complements as we assess the contribution of poor consumption in spell t to total, period-long poverty P_T . In other words, whenever we assess the extent of consumption shortfall in a particular spell, our valuation includes the memory of the shortfall if any in the last period. There is a close similarity to the literature on multidimensional poverty, where different attributes are assessed in relation to each other. Just as in multidimensional assessment, the fact that consumption (shortfalls) in any two spells must be combined into one composite leaves the gate open to questions on whether they complement (or substitute for) each other. In our case, complementarity is the only intuitive answer, since it is the fear that poor previous consumption may compound current hardship what motivates the *Prolonged poverty* axiom. The measure in (21) – one of many possibilities – allows for this complementarity, ensuring that poverty spells are valued higher in overall period poverty if they follow after another poverty spell.

7. Vulnerability in a dynamic world

The entire discussion thus far has considered poverty in a world with time, but with no risk. When constructing a measure of poverty over time *ex-post*, building on past observed outcomes in the standard of living, then this may be acceptable. Such a measure values actual realizations of a trajectory of the standard of living. However, using the earlier analysis when looking forward into the future to assess different paths of the standard of living, we implicitly assume perfect foresight: we know the realization of the standard of living without any uncertainty. In itself, such exercises are useful: for example, to compare trajectories under different policies or interventions. But one striking feature of such assessment is that it is unlikely to be done in a world of certainty, and risk should feature.

In Calvo and Dercon (2006), a measure of vulnerability as the “threat of poverty” has been derived. In particular, a set of desiderata has been proposed, borrowing from the standard poverty literature, and incorporating axioms that capture desirable properties stemming from the need to aggregate over states of the world. In the annex, an extract of this paper is given. The intuition is to provide an aggregate over some transformation of outcomes in all states of the world, whereby outcomes in each state are assessed relative to the poverty line. This gives a metric of the threat of poverty, before uncertainty has been resolved, and not of poverty itself. As appendix 1 shows, the desiderata include a focus axiom, symmetry over states, continuity and differentiability, scale invariance, normalisation, probability-dependent effect of outcomes, a probability-transfer axiom between states, and risk sensitivity (so increased risk raises vulnerability). If we impose an assumption of constant relative risk sensitivity, then it is shown that the preferred vulnerability measure will be the expected value of the Chakravarty measure of poverty:

$$V_{(\alpha)} = 1 - E\left(\frac{\tilde{y}_t}{z}\right)^\alpha, \text{ with } 0 < \alpha < 1. \quad (22)$$

E is the expected value operator, and α regulates the strength of risk sensitivity – as α rises to 1, we approach risk-neutrality. It is crucial to note that as defined by (22), vulnerability becomes greater whenever uncertainty rises, even if all in all expected outcomes remain unaltered. Thus, a normative choice is made to ensure that risk *per se* is bad and compounds expected hardship.

Note also that \tilde{y}_t is a vector consisting of \tilde{y}_{it} , censored outcomes for each state of the world i at time t . Although forward-looking, this measure is still essentially timeless: possible outcomes in timeless states are considered before the veil of uncertainty is lifted and before a particular state has been realized. Nevertheless, its desirable properties when constructing a measure of the threat of poverty mean that it could be used as a candidate for period-by-period outcomes before aggregation in an intertemporal measure of poverty. In particular, consider an amended version of (8), which in itself was based on the Chakravarty measure of poverty:

$$V_T(y_1, y_2, \dots, y_T) = \sum_{t=1}^T \beta^{T-t} \left[1 - E \left(\frac{y_t}{z} \right)^\alpha \right], \text{ with } 1 < \alpha \text{ and } 0 < \beta \quad (23)$$

This can be considered a forward-looking and dynamic measure of vulnerability, consistent with a form of 'FTAA'-case, where allowing for uncertainty requires that the final stage (after focus in each state and transformation) is extended to a double exercise: aggregation first takes place over all states of the world in each t , and then it operates over all periods of time. In each period, it satisfies a set desiderata that appear reasonable when assessing poverty *ex-ante* in risky world, as a metric of the threat of poverty. Even if the presence of risk will affect the exact formulation of the intertemporal desiderata, it appears clear that versions of the intertemporal axioms related to *Monotonicity*, *Increasing cost of transfers*, and *Sub-period decomposability* apply as well.

In other words, we have a measure of forward-looking intertemporal poverty, as a measure of the extent of the threat of poverty in the future, providing a clear ordering of different possible trajectories for individuals. It comes closer than any of its predecessors to providing a direct measure of 'chronic' poverty, in that it does not just assesses poverty in one period, nor assesses poverty in a risk-free world. It offers an exact way of ordering very different and complex trajectories, including the threat of poverty and deprivation implied *ex-ante* for those whose trajectory in expectation contains serious spells of severe deprivation, even if *ex-post* they do not always become realized.

8. An example from Ethiopia

To illustrate the insights that can be gained from a variety of measures of intertemporal poverty, we use data from rural Ethiopia. The Ethiopian Rural Household Survey has collected data on about 1450 households over the course of 10 years, in the form of six unequally spaced rounds. Here we drop round, which was collected in the second half of 2004, as it was collected in a distinctly different season and only about 6 months after the first round of 2004. The result is data from 1994, 1995, 1997, 1999 and 2004. We use data on consumption per capita, deflated to be

expressed in 1994 prices. The consumption aggregate is based on careful recording of consumption from own production, purchased items and gifts, and is predominantly food, at about 75 percent reflecting the relative poverty of households in rural Ethiopia. The data is relatively highly clustered, from only 15 communities, but reasonably well spread across the country. Round-by-round attrition was low, although we focus in the rest of the analysis on 1187 observations with complete information in all rounds. More details can be found in Dercon and Krishnan (2000). Using a poverty line not dissimilar from the national poverty line, at about 8.50 US dollars birr per capita per month, we find that the head count of poverty declined in this period, from 48 percent in 1994 and even 55 percent in 1995, to 33 percent by 1997 (an exceptionally good harvest year) and 36 and 35 percent in respectively 1999 and 2004. Still, there is considerable churning, and combined with the gradual decreasing poverty levels and possibly some problems of measurement error, we find that that only 18 percent of the households were never poor and 7 percent were poor in all rounds.

Table 1 **Poverty episodes 1994 to 2004 (based on 5 rounds)**

	Percentage of households (1)
Never Poor	18
Poor once	22
Poor in 2 out of 5 rounds	23
Poor in 3 out of 5 rounds	16
Poor in 4 out of 5 rounds	14
Poor in all rounds	7

Source: Ethiopia Rural Household Survey (based on 1187 observation with data in all 5 rounds).

Using these data, we calculated a number of different poverty measures summarizing these poverty experiences, using 1187 observations . First, and for comparison we calculated the squared poverty gap (the Foster-Greer-Thorbecke measure with $\alpha=2$) in the base year, 1994 and the final year, 2004. We find that it almost halved from 0.120 to about 0.065. In terms of intertemporal measures, we calculated measure (7) with $\alpha=2$, an FGT-style measure in which the focus axiom is applied before transformation and aggregation, so that no compensation is allowed between periods. We also use the assumption of equal-valued spells, i.e. $\beta=1$. Although measure (7) is not scaled by the

number of periods, dividing it by 5 gives a direct way to compare it with the period-by-period squared poverty gaps. Its scaled mean value is 0.089 is consistent with the nature of the decline in poverty in this period. Next, we calculated measure (12), effectively the Jalan and Ravallion (2000) measure, a squared poverty gap measure based on mean consumption in this period, allowing for compensation and equal-valued spells (with $\alpha=2$ and $\beta=1$). Its mean value of 0.025 suggests how strong the impact is of allowing for compensation, i.e. for aggregation before the focus axiom is applied. Further, we calculated two indexes of poverty, based on (7) but relax the assumption of equal-valued spells, by focusing on an index that values more recent years less than the past ($\beta=0.85$) and an index that values the present more than the past ($\beta=1.15$). Finally, we introduced sequence-sensitivity, using measure (22), which values poverty gaps only to the extent that one was poor in the previous year, using $\rho=0.90$, nesting it with the other cases by choosing $\alpha=2$ and $\beta=1$. The actual values of these last three indexes cannot quite be compared with the other indexes shown.

For empirical relevance, we need to ask whether these different measures of poverty tell us any different messages about poverty. As these measures are different non-linear transformations of underlying consumption measures, a first appropriate way to compare these measures would be to look at rank correlations: do they order people differently? Table 2 gives Spearman correlation coefficients for all these measures.

Table 2: Spearman Rank Correlation between Different Poverty Measures

	Sq Pov Gap 1994	Sq Pov gap 2004	FTA (7)	AFT (12)	Seq FTA (22)	FTA (7), ($\beta=0.85$)	FTA (7), ($\beta=1.15$)
Sq Pov Gap 1994 ($\alpha=2$)	1						
Sq Pov gap 2004 ($\alpha=2$)	0.166	1					
FTA (7), ($\beta=1$, $\alpha=2$)	0.690	0.462	1				
AFT (12) ($\beta=1$, $\alpha=2$)	0.553	0.468	0.689	1			
Seq FTA (22) ($\beta=1$, $\alpha=2$, $\rho=0.90$)	0.662	0.371	0.824	0.715	1		
FTA (7), ($\beta=0.85$, $\alpha=2$)	0.751	0.404	0.993	0.678	0.821	1	
FTA (7), ($\beta=1.15$, $\alpha=2$)	0.633	0.516	0.994	0.690	0.813	0.974	1

Source: calculated from the Ethiopian Rural Household Survey by authors.

As could be expected, all measures are positively (significantly) correlated, but some interesting differences emerge. Poverty in 1994 and in 2004 is relatively weakly correlated, partly reflecting the overall decline. Among the intertemporal measures,

using the FTA (7) measure with different discount rates does not appear to matter $\alpha\beta$ much for the ranking of households, with high correlations with each other. Choices on the sequence of focus, transformation and aggregation appears to matter most, with a correlation of about 0.69 between the AFT (12) and the FTA (7) measures with otherwise equal values for α and β . Adjusting for the sequence of poverty outcomes matters, but the correlation remains high with the other AFT measures. At least in these data, choices on allowing for compensation appear to most important, while cross-section poverty estimates for a population may give the wrong impression on intertemporal poverty outcomes and rankings.

Of course, much of this difference may be due to a different treatment of measurement error in welfare outcomes, entailed by each of these intertemporal poverty measures. More in general, the differences in poverty may be due to individual specific attributes hardly observable to researcher. One way of assessing whether our interpretation on the nature of poverty is different across measures is by constructing a ‘poverty profile’, a multivariate description of the correlates of poverty in these data, effectively are whether we identify different types of households to be poor using these different measures. This definitely not an exploration of a causal relationship between any of the factors identified and poverty – more careful analysis would be required – but it can give some sense of whether different concepts of intertemporal poverty result in different implications, for example, when trying to target poor population on the basis of generic characteristics. Table 3 gives the correlates of some of the different poverty measures used in table 2: the poverty gap in 1994, the FTA (7), the AFT (12) and the sequential FTA. The last two FTA measures, with different discount rate, were not used as they are very highly correlated with the FTA (7). As the poverty measures used are all censored, we use a tobit model with censoring at zero. Table 3 reports the coefficients.

Table 3 Correlates of Poverty Measures (Tobit model)

	Poverty Gap 1994	AFT (7)	FTA (12)	Seq AFT (7)
head at least primary ed.	-0.094 [2.35]**	-0.054 [4.66]***	-0.088 [3.36]***	-0.025 [3.56]***
head some primary	-0.075 [2.68]***	-0.02 [2.39]**	-0.013 [0.77]	-0.006 [1.29]
ln land per capita (ha)	-0.046 [3.48]***	-0.019 [4.89]***	-0.03 [3.64]***	-0.008 [3.42]***
sex of the head is male	-0.068 [2.49]**	-0.007 [0.83]	0.001 [0.04]	-0.005 [0.97]
no of female adults	0.029 [3.13]***	0.01 [3.32]***	0.013 [2.24]**	0.006 [3.95]***
no of girls 5-15	0.026 [2.69]***	0.005 [1.84]*	0.008 [1.37]	0.004 [2.61]***
no of girls 0-5	0.031 [2.18]**	0.021 [4.83]***	0.041 [4.74]***	0.011 [4.42]***
no of females 65+	0.037 [1.17]	0.004 [0.46]	-0.011 [0.50]	0.009 [1.59]
no of male adults	0.015 [1.61]	0 [0.12]	-0.007 [1.11]	0.001 [0.63]
no of boys 5-15	0.039 [4.14]***	0.014 [4.85]***	0.02 [3.40]***	0.006 [3.67]***
no of boys 0-5	0.056 [3.80]***	0.022 [4.88]***	0.038 [4.26]***	0.011 [4.45]***
no of males 65+	-0.047 [0.96]	-0.006 [0.43]	-0.024 [0.76]	-0.008 [0.93]
distance to town (km)	0.021 [10.60]***	0.006 [10.43]***	0.013 [9.04]***	0.004 [10.84]***
coeff. variation rainfall	0.006 [5.69]***	0.002 [4.82]***	0.003 [3.86]***	0.001 [3.45]***
is road accessible trucks	-0.22 [7.56]***	-0.072 [8.66]***	-0.164 [7.42]***	-0.04 [7.75]***
village mean land p.c.	0.228 [7.48]***	0.051 [5.64]***	0.066 [3.55]***	0.026 [5.03]***
village mean male adults	-0.049 [0.71]	-0.039 [1.87]*	0.052 [1.17]	-0.024 [1.98]**
village mean fem adults	0.397 [6.29]***	0.188 [9.93]***	0.178 [4.58]***	0.095 [8.79]***
Constant	-0.619 [7.90]***	-0.212 [9.14]***	-0.597 [10.33]***	-0.152 [10.99]***
Observations	1125	1125	1125	1125

Absolute value of t statistics in brackets

* significant at 10%; ** significant at 5%; *** significant at 1%

The correlates used include educational characteristics of the head (whether completed primary education or more, and whether some primary education, with the base group no education), land holding in hectares and per capita, the sex of the head, demographic composition of the household (number of male and female adults, children and elderly), and a number of village characteristics: the distance to the

nearest town in kilometers, whether there is a road passing the village that is accessible to trucks, buses and cars, and the coefficient of variation of rainfall in the village and finally, a few mean village characteristics, such as the mean land holding per capita, and the mean number of female and male adults per household (as there are substantial differences in land holdings and in demographic composition across villages).

The most striking insight from the table is that the differences between the different intertemporal measures of poverty appears relatively small: in any case, in terms of significance, the same variables appear to stand out, with the expected signs: education, land, distance to towns, road access and weather variability. Demographic characteristics also matter but not the sex of the head. Strikingly, even the profile based on the 1994 squared poverty gap offers broadly a similar set of correlates. Obviously, this does not mean that the same people are being predicted as being poor across equations.

It is difficult to interpret the differences in the size of the coefficients across equations, as the left hand side variables are rather different and most are not directly comparable. To highlight better the different interpretations across the regressions, we can compare the marginal effects *relative* to the mean of each left hand size variable. In other words, we can establish the percentage change on each poverty measure from a change in one of the explanatory variables. The relevant marginal effects are not the coefficients given in table 3 as the zeros in the data can be given direct meaning (a zero squared poverty gap is a zero squared poverty gap, and not some unobserved negative poverty). The coefficients in table 3 give the marginal effects relative to the underlying latent variable of the statistical model which is assumed to take on negative values. Instead, we use marginal effects based on the unconditional expected value, evaluated at the mean of all explanatory variables. Expressing these as a percentage of the mean dependent variable for each poverty measure, we obtain table 4.

Table 4 Percentage change in Poverty Index from Marginal Change in Characteristics

	AFT (7)	FTA (12)	Seq AFT (7)
From no education to primary completed	-0.43	-0.39	-0.48
Doubling land per capita	-0.17	-0.19	-0.18
Reducing distance to town by one kilometer	-0.05	-0.08	-0.09
From bad or no road to road accessible for trucks/bus	-0.63	-1.01	-0.91

Source: calculated from results in table 3

These results are suggestive, as there are some interesting differences in the order of magnitudes of the relative marginal effects. The most striking differences relate to the infrastructure variables: using the FTA (12) measure (i.e. allowing for compensation over time) suggests that living nearer to towns or with better roads is associated with considerably lower poverty than implied by the AFT (7). Education improvements are more strongly related to the AFT measures, especially the measure that effectively only counts repeated poverty episodes. In short, when using poverty measures over time, the way aggregation over time is done will affect the characteristics that will be especially highlighted in poverty profiles as correlated with lower poverty.⁴

Conclusions

This paper has offered a discussion of a number of issues related to measuring poverty over time. It has highlighted some of the key normative decisions that have to be taken. In particular, we have highlighted the role of compensation over time (whether poverty spells can be compensated for by non-poverty spells); the issue of the discount rate (whether each spell should be given an equal weight); and the issue of the role of persistence (whether repeated spells should be given a higher weight). We have offered a number of plausible poverty measures, each with different assumptions regarding these key issues. We have also shown how these insights can be used to construct a forward looking measure of vulnerability. Applying a number of these measures to data from rural Ethiopia, it is shown that while correlations are high, there would still be considerable differences in ranking households by poverty

⁴ As is well known with poverty profiles, these results have only limited policy implications, as these correlates are not shown to be causal factors, and even if they were, the relative cost of intervening in terms of infrastructure, land or education would have to be taken into account.

according to different measures, especially those that have different views on the role of compensation. Turning to a multivariate poverty profile, it was shown that while similar factors are significant, their relative importance in identifying intertemporal poverty is different according to the measure used to summarize poverty.

Appendix 1

A family of individual vulnerability measures (based on Calvo and Dercon (2006))

Let individual vulnerability (V) be measured by $V=v(z,\mathbf{p},\mathbf{y})$, where z is the poverty line, and \mathbf{p} and \mathbf{y} are k -dimensional vectors, containing state-of-the-world probabilities and outcomes, respectively – i.e., p_i is the probability of the i -th state occurring, with outcome y_i . We impose $y_i \geq 0$. It may be easiest to think of these outcomes as consumption levels in each possible state of the world, especially if poverty is defined as usual as a shortfall in consumption. We remark that we mean outcomes *after all consumption-smoothing efforts have been deployed*. In other words, their variability across states is taken as a final word, with no scope for reducing it further, e.g. by formal insurance, risk-sharing, or precautionary savings.

For each state, define ‘censored outcome’ \tilde{y}_i by $\tilde{y}_i \equiv \text{Min}(y_i, z)$, and the ‘rate of coverage of basic needs’ x_i by $x_i \equiv \tilde{y}_i/z$, so that $0 \leq x_i \leq 1$. Vectors $\tilde{\mathbf{y}}$ and \mathbf{x} are defined correspondingly. \mathbf{e}_i stands for a k -dimensional vector whose elements are 0, except for the i -th one, which equals 1. We close our notation with vectors $\hat{\mathbf{y}}$ and $\tilde{\mathbf{y}}^c$. Their elements are all equal to \hat{y} and \tilde{y}^c , respectively, which in turn are defined by $\hat{y} = \sum_{i=1}^k p_i \tilde{y}_i$ and $v(z,\mathbf{p},\tilde{\mathbf{y}}) = v(z,\mathbf{p},\tilde{\mathbf{y}}^c)$. Note that \tilde{y}^c can be written as a function $\tilde{y}^c(z,\mathbf{p},\tilde{\mathbf{y}})$ and will shortly be called the risk-free equivalent to the set of prospects described by $(z,\mathbf{p},\mathbf{y})$, in the sense that it yields the same degree of vulnerability. \hat{y} is the expected value of \tilde{y}_i .

We propose eight desiderata. The first is the FOCUS AXIOM, which imposes $v(z,\mathbf{p},\mathbf{y}) = v(z,\mathbf{p},\tilde{\mathbf{y}})$. Our measure will thus disregard outcome changes above the poverty line. If vulnerability is understood as a burden caused by the threat of future poverty, it should not be compensated by simultaneous (ex-ante) possibilities of being well-off. In consequence, high vulnerability is not necessarily tantamount for grim overall expected wellbeing (as arguably in Ligon and Schechter), since the ‘promise’ of richness in some states can raise welfare expectations, with no bearing on vulnerability.

Imagine that a farmer faces two scenarios: rain (no poverty) or drought (poverty). Does she become less vulnerable if the harvest in the rainy scenario improves? Our answer is ‘no’. *Poverty is as bad a threat as before*. It is as likely as before, and it is potentially as severe as before.

According to this axiom, ‘excess’ outcomes $y_i - z > 0$ are ‘wasteful’ and can be ignored, as far as vulnerability is concerned. Taking this for granted, the remaining axioms can be presented as follows:

SYMMETRY OVER STATES: $v(z, \mathbf{p}, \tilde{\mathbf{y}}) = v(z, \mathbf{Bp}, \mathbf{B}\tilde{\mathbf{y}})$, where \mathbf{B} is any $k \times k$ permutation matrix. All states receive the same treatment, and the only relevant difference between two states of the world i and j is the difference in their outcomes (y_i, y_j) and probabilities (p_i, p_j).

CONTINUITY AND DIFFERENTIABILITY. Function $v(z, \mathbf{p}, \tilde{\mathbf{y}})$ is continuous and twice-differentiable in \mathbf{y} , for tractability and to preclude abrupt reactions to small changes in outcomes.

SCALE INVARIANCE. $v(z, \mathbf{p}, \tilde{\mathbf{y}}) = v(\lambda z, \mathbf{p}, \lambda \tilde{\mathbf{y}})$ for any $\lambda > 0$. Our measure will not depend on the unit of measure of outcomes.

NORMALISATION. $\text{Min}_{\tilde{\mathbf{y}}} [v(z, \mathbf{p}, \tilde{\mathbf{y}})] = 0$ and $\text{Max}_{\tilde{\mathbf{y}}} [v(z, \mathbf{p}, \tilde{\mathbf{y}})] = 1$. We impose closed boundaries to facilitate interpretation and comparability.

PROBABILITY-DEPENDENT EFFECT OF OUTCOMES. For $-c < \tilde{y}_i < z$ and $p_i p'_i \neq 0$, $v(z, \mathbf{p}, \tilde{\mathbf{y}}) - v(z, \mathbf{p}, \tilde{\mathbf{y}} + c\mathbf{e}_i) = v(z, \mathbf{p}', \tilde{\mathbf{y}}') - v(z, \mathbf{p}', \tilde{\mathbf{y}}' + c\mathbf{e}_i)$ if and only if $p_i = p'_i$ and $\tilde{y}_i = \tilde{y}'_i$. Should \tilde{y}_i change, the consequent effect on vulnerability is not allowed to depend on the outcomes or probabilities of other states of the world – for a given p_i , the change in vulnerability depends only on \tilde{y}_i .⁵ In the opposite direction, the effect must be sensitive to the likelihood of that particular state of the world. Note that $p_i p'_i \neq 0$ discards ‘impossible’ states ($p_i = p'_i = 0$).

PROBABILITY TRANSFER. For every $p_j \geq d > 0$, $v(z, \mathbf{p} + d(\mathbf{e}_i - \mathbf{e}_j), \tilde{\mathbf{y}}) \begin{cases} \leq \\ \geq \end{cases} v(z, \mathbf{p}, \tilde{\mathbf{y}})$ if $\tilde{y}_i \begin{cases} \geq \\ \leq \end{cases} \tilde{y}_j$. If

\tilde{y}_i is greater than or at least equal to \tilde{y}_j , then vulnerability cannot increase as a result of a probability transfer from state j to state i . Likewise, if \tilde{y}_i is lower than or at most equal to \tilde{y}_j , then vulnerability cannot decrease. Going back to the example of the farmer facing rain and drought, we say that she becomes more vulnerable if a drought becomes more likely, at the expense of the rainy scenario (or at least, her vulnerability does not lessen as a result).

RISK SENSITIVITY. $v(z, \mathbf{p}, \tilde{\mathbf{y}}) > v(z, \mathbf{p}, \hat{\mathbf{y}})$. Vulnerability would be lower if the expected (censored) outcome \hat{y} were attained in all states of the world and uncertainty were thus removed. In other words, greater risk raises vulnerability.⁶ Thus we link up with our first intuition about vulnerability, as a concept aiming to capture the burden of insecurity, the fact that hardship is also related to fear of future threats.

Alternatively, resorting to the risk-free equivalent \tilde{y}^c , the same axiom could be expressed as $\tilde{y}^c / \hat{y} < 1$. Expected outcome is unevenly and ‘inefficiently’ spread across states of the world, in the sense that a similarly low degree of vulnerability would result from $\tilde{y}^c < \hat{y}$ being secured in every state. \tilde{y}^c / \hat{y} reflects this ‘efficiency loss’.

⁵ A possible counterargument could run ‘in fact, there could be some relief in considering that one could have done much better had the odds been more fortunate’ (or to the contrary, ‘she may rue having missed a better possible outcome, with no fault on her part, and thus her misery will be greater’). We ignore such counterarguments for the sake of tractability. In doing so, we simply adhere to the common concept of poverty as mere failure to reach a poverty line, with no regard for ‘subjective’ subtleties.

⁶ We implicitly define the increase in risk as a probability transfer ‘from the middle to the tails’, in keeping with one of the Rothschild-Stiglitz senses of risk.

CONSTANT RELATIVE RISK SENSITIVITY. For $\kappa > 0$, $\kappa \tilde{y}^c(z, \mathbf{p}, \tilde{\mathbf{y}}) = \tilde{y}^c(z, \mathbf{p}, \kappa \tilde{\mathbf{y}})$. A proportional increase by κ in the outcomes of all possible states of the world leads to a similar proportional increase in the risk-free equivalent \tilde{y}^c . While risk sensitivity ensures $\tilde{y}^c/\tilde{y} < 1$, we now require this ratio (or ‘efficiency loss’) to remain constant if all state-specific outcomes increase proportionally.

As compared to the previous axioms, this final property seems less compelling. Still, we find it attractive for its contribution both to narrowing down the families of acceptable measures to only one, and to securing that risk sensitivities receive an appropriate treatment. As for this second point, Ligon and Schechter (2003) were the first to point out that some existing vulnerability measures hid some awkward assumptions, e.g. risk sensitivity increasing in initial income, at odds with most empirical findings on risk attitudes (e.g. Binswanger 1981).

Needless to say, we are avoiding here terms such as ‘risk aversion’ or ‘utility’. We intend our choice of language to convey our view of vulnerability as distinct from expected utility, if only to stress our departure from proposals where vulnerability boils down to some form of bad ‘overall’ expectations (e.g. Ligon and Schechter). On the other hand, parallels should be obvious. In fact, the proof of the following theorem heavily draws on results from expected utility theory (mainly Pratt 1964), necessarily with some departures due to the specific traits of our vulnerability concept. For this reason and for brevity, it is not provided, but it is available on request.

THEOREM 1 – If all the axioms above are satisfied, then

$$V_{(\alpha)} = 1 - E[x^\alpha], \text{ with } 0 < \alpha < 1. \quad (1)$$

E is the expected value operator, and we recall $x_i \equiv \tilde{y}_i/z$ is the rate of coverage of basic needs, and $0 \leq x_i \leq 1$. We highlight the simplicity of this single-parameter family of measures $V_{(\alpha)}$.⁷ Of course, α regulates the strength of risk sensitivity – as α rises to 1, we approach risk-neutrality.

A few remarks are in place. First, for those facing no uncertainty and with known $x_i = x^* < 1$ for all i , $V_{(\alpha)} > 0$. If vulnerability is about the threat of poverty, certainty of being poor is but a dominant, irresistible threat. The concept is not confined to those whom the winds might blow into poverty or out from it. Vulnerability is about risk, but not only about it.

Second, it is easy to prove that $V_{(\alpha)}$ is equal to the probability of being poor only if outcomes are expected to be zero in every state of the world where the individual is poor. If vulnerability were measured as expected FGT₀ (as in Chaudhuri et al. 2002), then vulnerability would be overestimated. Ligon and Schechter have pointed out the shortcomings of other FGT choices.⁸

⁷ For instance, if our last axiom (constant relative risk sensitivity) were replaced by constant absolute risk sensitivity [$\kappa + \tilde{y}^c(z, \mathbf{p}, \tilde{\mathbf{y}}) = \tilde{y}^c(z, \mathbf{p}, \tilde{\mathbf{y}} + \kappa)$, for $\kappa > 0$], the less attractive measure $V_{(\beta)} = 1 - E[\{e^{\beta(1-x)} - 1\} / \{e^\beta - 1\}]$, with $\beta > 0$, would result.

⁸ More precisely, we should speak about expected individual poverty, as measured by the function implicit in the corresponding aggregate FGT index, as in Foster, Greer and Thorbecke (1984).

Finally, $V_{(a)}$ can still be assimilated into the expected-poverty approach to vulnerability, provided poverty is measured as in Chakravarty (1983). In some sense, one of the contributions of this paper is to identify the Chakravarty poverty index as the best choice if the poverty analysis moves from static poverty on to vulnerability.

References

- Anand, S. and K. Hanson, 1997, Disability-Adjusted Life Years: A Critical Review, *Journal of Health Economics*, 16, 685-702.
- Baulch. B. and J. Hoddinott, 2000, "Economic Mobility and Poverty Dynamics in Developing Countries", introduction to special issue, *Journal of Development Studies*.
- Binswanger, H. P, (1981). "Attitudes toward Risk: Theoretical Implications of an Experiment in Rural India," *Economic Journal*, Royal Economic Society, 91(364): 867-90.
- Bourguignon, F. and S.R. Chakravarty, 2003, The Measurement of Multidimensional Poverty, *Journal of Economic Inequality*, vol. 1, no. 2, 1569-1721
- Calvo, C. and S. Dercon, 2006, *Vulnerability to Poverty*, Oxford University, mimeo.
- Chakravarty, S.R., 1983, "A New Index of Poverty" *Mathematical Social Sciences*, vol. 6, 307-313.
- Chaudhuri, S., and J. Jalan (2002). Assessing Household Vulnerability to Poverty from Cross-Sectional Data: A Methodology and Estimates from Indonesia. Columbia University Discussion Paper 0102-52.
- Dercon, S., and P. Krishnan. 2000. Vulnerability, poverty and seasonality in Ethiopia. *Journal of Development Studies* 36 (6): 25-53.
- Foster, J., J. Greer, and E. Thorbecke, 1984. "A Class of Decomposable Poverty Measures." *Econometrica* 52(3): 761-766.
- Foster, J. and T. Shorrocks 1991, Sub-group Consistent Poverty Indices, 59 (3), 689-709.
- Foster, J., 2007. A Class of Chronic Poverty Measures, mimeo.
- Jalan, J. and M. Ravallion, 2000, 'Is Transient Poverty Different? Evidence from Rural China', *Journal of Development Studies*, Vol.36, No.6.
- Kanbur, R. and D. Mukherjee 2006. "Premature Mortality and Poverty Measurement", Mimeo, Cornell University.
- Ligon, E. and L. Schechter (2003). "Measuring Vulnerability." *The Economic Journal* 113(486): C95-C102.
- Pratt, J. W. (1964). "Risk Aversion in the Small and in the Large." *Econometrica* 32(1/2): 122-136.
- Tsui, K. 2002. "Multidimensional Poverty Indices." *Social Choice and Welfare* 19(1):69-93.