

A CLASS OF CHRONIC POVERTY MEASURES

by

James E. Foster^{*}

Draft 4

November 29, 2006

Work in progress: please do not quote.

^{*}Department of Economics, Vanderbilt University. This paper was written for the CPRC Workshop on Concepts and Methods for Analysing Poverty Dynamics and Chronic Poverty, Manchester, UK, 23-25 October 2006. The author would like to thank Maria Emma Santos for expert research assistance and the participants of the CPRC Workshop and the LACEA/LAMES Conference in Mexico City for helpful comments.

ABSTRACT

This paper presents a new family of chronic poverty measures based on the P_α poverty measures of Foster, Greer, and Thorbecke (1984). The chronically poor are identified using two cutoffs: a standard poverty line, which identifies the time periods during which a person is poor; and a duration cutoff, which is the minimum percentage of time a person must be in poverty in order to be chronically poor. The new family of chronic poverty measures is constructed by raising the (per-period) normalized gaps of the chronically poor to a power $\alpha \geq 0$ and then aggregating. The resulting indices, which can be viewed as duration adjusted P_α measures, satisfy a battery of properties for chronic poverty indices, including time monotonicity and population decomposability. An illustrative application of the family is provided using data from Argentina.

CONTENTS

- I. INTRODUCTION**
- II. TRADITIONAL POVERTY MEASUREMENT**
- III. THE MEASUREMENT OF CHRONIC POVERTY**
 - A. NOTATION*
 - B. IDENTIFYING THE CHRONICALLY POOR*
 - C. CHRONIC POVERTY AND AGGREGATION*
 - D. PROPERTIES FOR CHRONIC POVERTY MEASURES*
- IV. AN EMPIRICAL ILLUSTRATION**
- V. CONCLUSIONS**

I. INTRODUCTION

Traditional measures of poverty based on cross-sections of income (or consumption) data provide important information on the incidence of material poverty, its depth, and distribution across the poor. However, they have little to say about another important dimension of poverty: its duration. Empirical evidence suggests that increased time in poverty is associated with a wide range of detrimental outcomes, especially for children.¹ If so, then this would provide a strong rationale for using a methodology for evaluating chronic poverty that explicitly incorporates “time in poverty”. This paper presents a new class of chronic poverty measures that can account for duration in poverty as well as the traditional dimensions of incidence, depth and severity.

There are several methodologies available for measuring chronic poverty using panel data. Two broad categories may be discerned, each with its own distinctive strategy for identifying the chronically poor.² The *components* approach, exemplified by Jalan and Ravallion (1998), constructs an average or permanent component of income and identifies a chronically poor person as one for whom this component lies below an appropriate poverty line.³ Variations in incomes across periods are ignored by this identification process and by the subsequent aggregation step when the data are brought together into an overall measure. The components approach to chronic poverty measurement is not especially sensitive to the time a family spends in poverty and, hence, may not be the best framework for incorporating duration into poverty measurement.

A second approach to evaluating chronic poverty – called the *spells* approach – focuses directly on the period-by-period experiences of poor families, and especially on the time spent in poverty. The identification of the chronically poor typically relies on a duration cutoff as well as a poverty line: Gaiha and Deolalikar (1993), for example, takes the set of chronically poor to be all families that have incomes below the poverty line in at least five of the nine years of observations, hence have a duration cutoff of 5/9. As for the aggregation step, most proponents of the spells approach use a very simple index of chronic poverty based on the *number* of chronically poor.⁴ While the number (or

¹ For example, longer exposure to poverty is associated with: increased stunting, diminished cognitive abilities and increased behavioral problems for children (Brooks-Gunn and Duncan, 1997); worse health status for adults (McDonough and Bergland, 2003); lower levels of volunteerism when poor children become adults (Lichter, Shanahan, and Gardner, 1999); and an increased probability of staying poor (Bane and Ellwood, 1986; Stevens 1994). See also the conceptual discussions of Yaqub (2003) and Clark and Hulme (2005).

² This division is due to Yaqub (2000); see also McKay and Lawson (2002).

³ Examples of the components approach can be found in Duncan and Rodgers (1991), Rodgers and Rodgers (1993), Jalan and Ravallion (1998), and Dercon and Calvo (2006), among others. Alternatively, one can estimate the permanent component based on household characteristics: see (who?).

⁴ See for example *The Chronic Poverty Report 2004-05*, p. 9, which uses a headcount. Duncan, Coe, and Hill (1984) and Gaiha and Deolalikar (1993) use the headcount ratio.

percentage) of chronically poor is an important statistic to keep in mind, it is a rather crude indicator of overall chronic poverty. In particular, it ignores the time a chronically poor family spends in poverty and hence a *time monotonicity* property that is especially relevant in the present context. In addition, other key dimensions of poverty, namely its depth and distribution, are utterly ignored by the index.

The present paper adopts the general methodology of the spells approach. Two distinct cutoffs are used for identifying the chronically poor – one in income space (the usual poverty line $z > 0$) and another governing the percentage of time in poverty (the duration line $0 < \tau < 1$). In words, a family is considered to be chronically poor if the percentage of time it spends below the poverty line z is at least the duration cutoff τ . For the aggregation step, this paper presents a new class of chronic poverty measures based on the P_α family proposed by Foster, Greer, and Thorbecke (1984), appropriately adjusted to account for the duration of poverty. All of the measures satisfy time monotonicity and an array of basic axioms, while certain subfamilies satisfy the multiperiod analogs of (income) monotonicity and the transfer principle. Associated measures of *transient* poverty are defined to evaluate poverty that is shorter in duration. Each chronic poverty measure (and its transient dual) satisfies decomposability, thus allowing the consistent analysis of chronic poverty by population subgroup. In particular, profiles of chronic poverty can be constructed to understand the incidence, depth and severity of poverty in a way analogous to the standard static case.

The paper proceeds as follows. Section II provides a brief overview of poverty measurement in a static environment to help ground the discussion of chronic poverty measurement. Section III introduces time into the analysis. The identification and aggregation steps are specified and the new family of chronic poverty measures is defined. Several sets of axioms are presented and used to evaluate of the new class of measures. Section IV provides a brief application of the technology to data from Argentina, while Section V concludes.

II. TRADITIONAL POVERTY MEASUREMENT

Following Sen (1976), poverty measurement can be broken down into two conceptually distinct steps: first, the identification step, which defines the criteria for determining who is poor and who is not; and, second, the aggregation step, by which the data on the poor are brought together into an overall indicator of poverty. The identification step is typically accomplished by setting a cutoff in income space called the *poverty line* and evaluating whether a person's resources are sufficient to achieve this level. There are many bases for selecting poverty lines, with a major differences being the information that is used in the setting of the line and in how the line changes over time. *Subjective* poverty lines consider information from surveys that ask participants how much it takes to get along. *Relative* poverty lines depend on the income standard achieved by a given society; a common example sets the poverty line at 50% of the median income. *Absolute* poverty lines may be purely arbitrary (such as the \$1 or \$2 per day lines used in World Bank illustrations) or may be initially derived from consumption studies. Note that in principle each type of line can be located at the low end or the high end of conceivable

cutoffs (e.g., a relative line at 1% of the median and an absolute line at \$15 per day); consequently, the use of an absolute line does not identify a person as being “absolutely impoverished.” Instead, the term “absolute” typically refers to the fact that the poverty line is to remain fixed during the time frame under consideration. In contrast, a thoroughgoing relative (or subjective) approach will have a different poverty line at each point in time as the income standards (or norms) change. This paper assumes that an absolute poverty line has been selected and that it is applicable at all time periods under consideration.⁵

The aggregation step is typically accomplished by selecting a particular poverty index or measure. Each index provides a different method of combining the income data and the poverty line into an overall indicator of poverty, and is more formally a function from the set of income distribution and poverty line pairs into the real numbers. The simplest and most widely used measure is the *headcount ratio*, which is the percentage of the given population that is poor. It is sometimes helpful to view the headcount ratio as a specific population average; indeed, if every person identified as being poor is assigned a value of “1” while every person outside the set of poor is assigned a value of “0”, then the headcount ratio is simply the mean of the resulting “0-1” vector.

A second method of aggregation is given by the (*per capita*) *poverty gap*, which is the aggregate amount by which the poor fall short of the poverty line income, measured in poverty line units, and averaged across the entire population. It too can be seen as a population average, with those outside of the set of the poor being assigned a value of “0”. However, instead of assigning a “1” to all poor persons, they are now given their own *normalized shortfall* (or the difference between their income and the poverty line, divided by the poverty line itself) before taking the population average. In contrast to the “all or nothing” approach of the headcount ratio, the poverty gap measures an individual’s level of poverty by the normalized shortfall, and then views poverty as the average value of this shortfall across society. Consequently, it is sensitive to variations in the incomes of the poor and indeed registers an increase when the shortfall of a poor person rises (*ceteris paribus*).

A third method of aggregation suggested by Foster, Greer and Thorbecke (1984) proceeds as above for each person who is not poor, but now *transforms* the normalized shortfalls of the chronically poor by raising them to a nonnegative power α to obtain the associated P_α or FGT measure. This approach actually includes both of the foregoing measures: P_0 is the headcount ratio and P_1 is the poverty gap measure. The *squared gap* measure P_2 from this family takes the square of each normalized shortfall, which has the effect of diminishing the relative importance of very small shortfalls and augmenting the effect of larger shortfalls – and hence emphasizing the conditions of the poorest poor in society. The index is a simple average of the squared normalized shortfalls across the

⁵ One could also imagine alternative types of hybrid approaches to setting poverty lines across space and time. See Foster and Szekely (2006).

population. While there are several other poverty measures in common use,⁶ this paper will focus on the FGT class of measures in general and the three measures, P_0 , P_1 , and P_2 , in particular, in developing a new class of chronic poverty measures.

Every poverty index offers a different view of poverty. One of the ways of clarifying these differences is to identify the properties or axioms satisfied by them. Each property captures a basic requirement for an aggregation method, and usually defines a form of stylized change in the income distribution that is then required to have a particular impact on measured poverty. The *focus axiom* requires that any increase in income for those outside the set of the poor should not affect the measured level of poverty. In other words, a poverty measure should focus on the set of the poor and their incomes, and not on the incomes of those outside the set of the poor. The *anonymity* property requires a permutation of incomes (by which the same incomes wind up in different hands) to leave poverty unaffected. In particular, this ensures that the income variable has been appropriately adjusted so there are no remaining factors associated with an individual's identity that should be taken into account in the measurement process.

The *replication invariance* axiom specifies that the index must be independent of the population size in that a replication of a given distribution (in which each income has been "cloned" a specific number of times) has the same level of poverty as the original distribution. The result is a formula that measures poverty on a per capita basis. *Scale invariance* requires that if both the poverty line and every income are scaled up or down by the same factor, then the poverty level should be unchanged. This leads us to view poverty as being measured in poverty line units, and thus allows coherent comparisons across time and space where poverty lines and incomes are changing. In the present case, where the poverty line is fixed, this will not be an essential aspect of the measure; however, it will prove useful in understanding what each index is actually measuring. Note that each of the three indices we previously discussed satisfies all four of these basic properties.

Poverty indices commonly differ from one another in their treatment of elementary changes in the income distribution among the poor. The *monotonicity* axiom specifies that a decrease in the income of a poor person should lead to an increase in the measured level of poverty. This is a natural property for a poverty measure to satisfy; in particular, if it were strongly violated (with a decrement in a poor income being associated with a lower level of poverty) it would lead to very odd policy prescriptions indeed. The headcount ratio violates the axiom – but just barely – since an income decrement among the poor leaves the headcount ratio unchanged. The other two measures mentioned above are appropriately sensitive to the incomes of the poor and hence satisfy the monotonicity axiom.

The *transfer* axiom says that if one poor person gives a small amount of income to a richer poorer person then poverty should rise. This is perhaps a less fundamental

⁶ See for example Sen (1976), Clark, Hemming and Ulph (1981), Chakravarty (1983), or the surveys of Foster and Sen (1997) and Foster (2005).

property than monotonicity, but following Sen (1976) it has received a great deal of support. In any case, if an index were found to *decrease* when a regressive transfer among the poor occurred, this could lead to unintuitive policy prescriptions to help the less poor at the expense of the poorest. The headcount can be subject to this criticism when the transfer pushes the richer poor person across the line. The poverty gap just violates this axiom in that it ignores the impact of such a transfer among the poor (but agrees with it when the transfer is large enough to push the recipient above the poverty line). The P_2 index, which puts greater weight on poorer incomes, satisfies the transfer axiom; in other words, it is sensitive to the distribution of income among the poor.

Our final two properties relate the overall poverty level to the levels of poverty in population subgroups. *Subgroup consistency* requires that whenever poverty falls in a given subgroup and stays the same or falls in the remaining subgroup (with respective population sizes being fixed), overall poverty should also fall. This can be justified from a practical policy perspective, since if overall poverty could rise when subgroup poverty levels decrease, success in alleviating poverty at the local level could be seen as a failure overall. Subgroup consistency ensures that such counterintuitive situations simply cannot arise. The second property is a slightly stronger condition in that it specifies the precise formula for linking up subgroup and overall poverty. *Decomposability* requires overall poverty to be the weighted average of subgroup poverty levels, where the weights are the population shares of the respective subgroups. This property has been used extensively in the empirical literature, and each of the three poverty measures we discussed satisfies it as well as subgroup consistency.

This paper will present a battery of properties for chronic poverty measures analogous to the ones that have been presented above, and will propose several others that do not have static analogs. The next section presents the new class of chronic poverty measures and investigates the properties they satisfy.

III. THE MEASUREMENT OF CHRONIC POVERTY

A main premise of chronic poverty evaluation is that poverty repeated over time has a greater impact than poverty that does not recur. This section discusses how poverty measurement can be altered to take into account the additional dimension of time in poverty. The first part begins with some important definitions and notation.

A. NOTATION

The basic data are observations of the income (or consumption) variable for a set $\{1, \dots, N\}$ of individuals at several points in time.⁷ Let $y = (y_i^t)$ denote the matrix of (nonnegative) income observations over time, where the typical entry y_i^t is the income of individual $i = 1, 2, \dots, N$ in period $t = 1, 2, \dots, T$. We adopt the convention that y is a $T \times N$ matrix (having height T and length N), so that each column vector y_i lists individual i 's incomes over

⁷ In general, an income variable is any single-dimensional, cardinally meaningful indicator of wellbeing. In the present case where per-period values are not transferable, the most natural interpretation of this variable may be as consumption flows.

time, while each row vector y^t gives the distribution of income in period t . It will prove helpful to use the notation $|y| = \sum_t \sum_i y_i^t$ to denote the sum of all the entries in a given matrix, and to define an analogous notation for vectors (hence $|y_i| = \sum_t y_i^t$ is the sum of i 's incomes across all periods while $|y^t| = \sum_i y_i^t$ is the total income in period t). It is assumed that incomes have been appropriately transformed to account for variations across time and household configurations so that a common poverty line z can be used to establish who is poor in each period.⁸

It is sometimes useful to express the data in terms of (normalized) shortfalls rather than incomes. Let g be the associated matrix of normalized gaps, where the typical element g_i^t is zero when the income of person i in period t is z or higher, while $g_i^t = (z - y_i^t)/z$ otherwise. Clearly, g is a $T \times N$ matrix whose entries are nonnegative numbers less than or equal to one. When an entry g_i^t is equal to zero, this indicates that the person's income is at least as large as the poverty line and hence is not in poverty; when an entry is positive, this indicates that the person's income falls below the line, with g_i^t being a measure of the extent to which that person is poor.⁹ We can similarly define the matrix s of squared normalized shortfalls by squaring each entry of g ; i.e., the typical entry of s is $s_i^t = (g_i^t)^2$.

Counting-based approaches to evaluating poverty ignore the extent of the income gap and instead only take into account whether the gap is positive or zero. It is therefore helpful to create another matrix h by replacing all positive entries in g with the number "1". Thus the typical entry h_i^t of h is "0" when the income of person i in period t is *not* below z , and "1" when y_i^t is below z . One statistic of interest in the present context is the *duration* of person i 's poverty, or the fraction of time the person is observed to have an income below z . Denote this by d_i and note that it can be obtained by summing the entries in h_i (the i th column of h) and dividing by the number of periods; i.e., $d_i = |h_i|/T$. In essence the duration is analogous to a headcount ratio, but defined for a *given* person over time, not across *different* people within the same period of time.

This paper's approach to chronic poverty will be based on the percentage of time a person spends in poverty. Toward this end, it will be useful to derive matrices from g (and also from s and h) that ignore persons whose duration in poverty falls short of a given target $\tau > 0$. Let $g(\tau)$ (and $s(\tau)$ and $h(\tau)$) be the matrix obtained from g (respectively s and h) by replacing the i th column with a vector of zeroes when $d_i < \tau$. In other words, the typical entry of $g(\tau)$, namely $g_i^t(\tau)$, is defined by $g_i^t(\tau) = g_i^t$ for all i satisfying $d_i \geq \tau$ while $g_i^t(\tau) = 0$ for all i having $d_i < \tau$ (with the analogous definition holding for s and h).

As the duration target τ rises from 0 to 1 the number of nonzero entries in the associated matrix falls, reflecting the progressive censoring of data from persons who are not

⁸ This is a footnote describing that in practice the poverty line may be the deflating mechanism and refers to the Foster 1998 paper and the US experience.

⁹ In this paper a person with an income of z is not poor; the alternative assumption (that z is a poor income level) could be adopted with a slight change in notation.

meeting the poverty duration requirement. It is clear that the specification $\tau = 0$ does not alter the original matrices at all; consequently, $g(0) = g$, $s(0) = s$, and $h(0) = h$. At the other extreme, where $\tau = 1$, any person who was out of poverty for even a single observation would have a column of zeroes; in other words, $g(1)$, $s(1)$, and $h(1)$ consider a person who fell out of poverty for *one period* indistinguishable to a person who was *always* out of poverty.

As in the static case, the measurement of chronic poverty can be divided into an identification step and an aggregation step. There are many potential strategies for defining identifying the chronically poor, but all have the effect of selecting a set Z of chronically poor persons from $\{1, \dots, N\}$. The aggregation step takes the set Z as given and associates with the income matrix y an overall level $K(y; Z)$ of chronic poverty. The resulting functional relationship K is called an *index*, or *measure*, of chronic poverty.

B. IDENTIFYING THE CHRONICALLY POOR

What can panel data reveal that cross sectional observations cannot? By following the same persons over several periods, one can discern whether the poverty experienced by a person in a given period is an exceptional circumstance or the usual state of affairs. With panel data, there is not one but several income observations linked to each individual, and this in turn leads to a wide array of potential methods for deciding when a person is chronically poor. One approach employed by Jalan and Ravallion (1998) bases membership in Z on a single comparison between the poverty line z and a composite indicator of the resources an individual has available through time. The specific income standard employed by Jalan and Ravallion is $\mu(y_i) = |y_i|/T$, the average or mean income over time; hence their method identifies as chronically poor any person whose mean income is below the poverty line. As noted above, this approach is not particularly sensitive to the duration of poverty. Nonetheless, it may make good sense when incomes are perfectly transferable across time and, accordingly, consumption can be completely smoothed. Anyone with an average income below z would in the best case be poor for every period; while a person with a mean of z or above could be out of poverty in every period. However, the assumption of perfect transferability may be difficult to sustain, particularly for poorer individuals; and if per-period incomes are even slightly less than perfectly substitutable, this procedure could easily misidentify persons.¹⁰

At the other extreme is the “spells” approach to identifying the chronically poor, which bases membership in Z upon the frequency with which one’s income falls below the poverty line. So, for instance, one might require a person to be poor 50% of the time or more, before identifying the person as chronically poor. A higher cutoff (say 70% of the time or more) would likely lead to a smaller set of persons being identified as chronically

¹⁰ The case of imperfect substitutability is considered by Foster and Santos (2006). Notice that there is a relationship between this “components” approach to identification and the “spells” approach when the two extreme aggregates – the maximum income across the T periods and the minimum income across the T periods – are used. The first identifies as chronically poor only those who are poor in *all* periods; the second identifies as chronically poor someone who is poor in *some* period.

poor, while a lower cutoff (such as 30%) would likely expand the set. Note, though, that this approach also contains within it an implicit assumption – that there is *no* possibility of transferring income across periods. Indeed, it is not entirely clear why a person with a tremendous amount of income (or expenditure) in period one, who is just barely below the poverty line in the remaining periods would be considered chronically poor, as may be required under this approach. Nonetheless, if (1) the poverty line is considered to be a meaningful dividing line between poor and nonpoor and (2) the observed data on income (or consumption) in each period faithfully reflects the constraint facing the person in the given period, then identifying chronic poverty with sufficient time in poverty makes intuitive sense.¹¹

This paper uses a dual cutoff approach to identifying the chronically poor: The first cutoff is the *poverty* line $z > 0$ used in determining whether a person is poor in a given period; the second is the *duration* line τ that specifies the minimum fraction of time that must be spent in poverty in order for a person to be chronically poor. Given the income matrix y and the poverty line z , the matrix h depicts the poverty spells for each person, and this in turn yields d_i , the fraction of time that person i is observed to have an income below z . Then given τ , the set of chronically poor persons is defined to be $Z = \{i : d_i \geq \tau\}$, or the set of all persons in poverty at least τ share of the time. Since Z depends on z and τ , the poverty index can be written as a function $K(y; z, \tau)$ of the income matrix and the pair of parametric cutoffs. The next section constructs several useful functional forms for $K(y; z, \tau)$.

C. CHRONIC POVERTY AND AGGREGATION

The first question that is likely to arise in discussions of chronic poverty is: How many people in a given population are chronically poor? The answer comes in the form of the *headcount* $Q(y; z, \tau)$ defined as the number of persons in Z . This statistic is often highlighted in order to convey meaningful information about the magnitude of the problem; however, when making comparisons, especially across regions having different population sizes, the *headcount ratio* $H(y; z, \tau) = Q(y; z, \tau)/N$ is used, where N is the population size of y . The measure H focuses only on the frequency of chronic poverty in the population and ignores all other aspects of the problem, such as the average time the chronically poor are in poverty, or the average size of their normalized shortfalls.

An example will help illustrate these concepts. Consider the income matrix

$$y = \begin{bmatrix} 3 & 9 & 7 & 10 \\ 7 & 3 & 4 & 8 \\ 9 & 4 & 2 & 12 \\ 8 & 3 & 2 & 9 \end{bmatrix}$$

¹¹ This footnote discusses consumption vs consumption flow vs income as the basis of measurement.

where the poverty line is $z = 5$ and the duration line is $\tau = 0.70$. The associated matrices of normalized poverty gaps, g , and of poverty spells, h , are given by

$$g = \begin{bmatrix} 0.4 & 0 & 0 & 0 \\ 0 & 0.4 & 0.2 & 0 \\ 0 & 0.2 & 0.6 & 0 \\ 0 & 0.4 & 0.6 & 0 \end{bmatrix} \quad h = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \end{bmatrix}$$

Summing the entries of h vertically and dividing by $T = 4$ yields the duration vector $d = (d_1, d_2, d_3, d_4) = (0.25, 0.75, 0.75, 0)$ and hence we see that $Q(y; z, \tau) = 2$ and so $H(y; z, \tau) = 0.5$; in this population, half of the persons (namely numbers 2 and 3) are chronically poor.

Now consider a thought experiment in which person 3 in the above example receives an income of 3 rather than 7 in period 1, and hence the normalized gap in that period becomes 0.40 and the entry in h becomes one. Then person 3 would still be chronically poor, but the poverty duration would now be $d_3 = 1.0$ rather than 0.75. What would happen to H in this instance? Clearly, it would be *unchanged* even though a chronically poor person has experienced an increment in the time spent in poverty. In other words, H violates an intuitive time monotonicity axiom (defined rigorously below). It can be argued that, while H conveys meaningful information about one aspect of chronic poverty, and hence is a useful “partial index”, it is a bit too crude to be used as an overall measure.¹²

There is a very direct way of transforming H into an index that is sensitive to changes in the duration of poverty. Consider the matrix $h(\tau)$ defined above, which leaves a column unchanged if the person is chronically poor, and otherwise replaces the column with zeroes. Let $d_i(\tau) = |h_i(\tau)|/T$ denote the associated duration level of person i , so that $d_i(\tau) = d_i$ for each chronically poor person and $d_i(\tau) = 0$ otherwise. Then the *average duration* among the chronically poor is given by $D(\tau) = (d_1(\tau) + \dots + d_N(\tau))/Q$. This is a second partial index that conveys relevant information about chronic poverty, namely, the fraction of time the average chronically poor person spends in poverty. Combining the two partial measures yields an overall index that is sensitive to increments in the time a chronically poor person spends in poverty as well as to increases in the prevalence of chronic poverty in the population. Define the *duration adjusted headcount ratio* $K_0 = HD$ to be the product of the original headcount ratio H and the average duration D or, equivalently, $K_0 = (d_1(\tau) + \dots + d_N(\tau))/N$.

K_0 offers a different interpretation of our thought experiment than the one provided by H . Return to the original situation in which person 3 is not poor in period 1. For $\tau = 0.70$, the relevant $h(\tau)$ matrix is given by

¹² See the discussion of partial indices in Sen and Foster (1997).

$$h(\tau) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \end{bmatrix}$$

and the respective column averages are given by $d_1(\tau) = d_4(\tau) = 0$ and $d_2(\tau) = d_3(\tau) = 0.75$. The headcount ratio is $H = 0.50$ while the mean duration is $D = 0.75$ so that the duration adjusted headcount ratio K_0 is initially 0.375. Now when person 3's income in period 1 becomes 3 rather than 7, the fraction of time spent in poverty rises to 1 for that person, while the mean duration among all chronically poor rises to 0.875. Consequently, even though H is unchanged, K_0 rises to about 0.438, with this higher overall level of chronic poverty being due to person 3's increased time in poverty.

The above example also shows that $K_0 = \mu(h(\tau)) = |h(\tau)|/(TN)$; in words, K_0 is the mean of the entries in matrix $h(\tau)$ or, equivalently, the total number of periods in poverty experienced by the chronically poor as given by $|h(\tau)|$, divided by the total number of possible periods across all people, or TN . In the above example, it is easy to see that the mean of the 16 entries in $h(\tau)$ is $6/16 = 0.375$, and hence this is the duration adjusted headcount index K_0 . Notice that if a chronically poor person were to have an additional period in poverty, this would raise an entry in the matrix $h(\tau)$ from zero to one, thereby causing the average value K_0 to rise, as noted above. K_0 satisfies the time monotonicity axiom defined in the next section.

There is no doubt that K_0 is less crude than H as an overall measure of chronic poverty. However, it still is remarkably insensitive to the actual conditions of the chronically poor. The matrix $h(\tau)$, upon which K_0 is based, is unaffected by changes in incomes (or normalized gaps) that preserve the signs the entries of $g(\tau)$, even if the magnitudes of the entries in $g(\tau)$ change dramatically. For example, if the income of person 3 in period 2 were decreased from 4 to 2, so that the normalized gap g_3^2 would rise from 0.2 to 0.6, the corresponding entry in h would obviously be unchanged (namely, $h_3^2 = 1$), and hence K_0 would remain the same. So a chronically poor person is now much poorer in period 2, and yet this fact goes unnoticed by the duration adjusted headcount measure. This is a violation of the (income) monotonicity axiom (defined rigorously in the next section).

What is missing from this measure is information on the *magnitudes* of the normalized gaps. Consider the matrix $g(\tau)$ defined above whose nonzero entries are the normalized gaps of the chronically poor. The number of nonzero entries in $g(\tau)$ – and hence $h(\tau)$ – is $|h(\tau)|$, while the sum of the nonzero entries in $g(\tau)$ is $|g(\tau)|$. The ratio $|g(\tau)|/|h(\tau)|$ indicates the average size of the normalized gaps across all periods in which the chronically poor are in poverty. The resulting *average gap* $G(\tau) = |g(\tau)|/|h(\tau)|$ provides exactly the type of information that would usefully supplement the adjusted headcount ratio. Define the *duration adjusted poverty gap index* $K_1 = K_0G$ to be the product of the adjusted headcount ratio K_0 and the average gap G or, equivalently, $K_1 = HDG$, the product of the three partial indices that respectively measure the prevalence, duration, and depth of chronic poverty.

This chronic poverty index provides a third perspective from which to view our numerical example. Given the duration cutoff $\tau = 0.70$, the matrix $g(\tau)$ associated with the original situation in which person 3 has an income of 4 in period 2 is given by

$$g(\tau) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0.4 & 0.2 & 0 \\ 0 & 0.2 & 0.6 & 0 \\ 0 & 0.4 & 0.6 & 0 \end{bmatrix}$$

The respective sum of entries is $|g(\tau)| = 2.4$ while the number of periods in poverty is $|h(\tau)| = 6$, and hence the average gap is $G = 0.40$. Given $H = 0.50$ and $D = 0.75$ from before, the resulting level of the duration adjusted poverty gap measure is $K_1 = 0.15$. Now suppose that the period 2 income of person 3 falls from 4 to 2. Clearly H and D are unaffected by this change, and so K_0 would likewise be unchanged. However, the average gap G would rise to about 0.47, and hence the duration adjusted gap would now be $K_1 = .175$, reflecting the worsened circumstances for person 3. K_1 rises as a result of the income decrement since it satisfies monotonicity.

The duration adjusted gap measure has a simple expression as the mean of the entries of the matrix $g(\tau)$, so that $K_1 = \mu(g(\tau)) = |g(\tau)|/(TN)$. In words, K_1 is the sum of the normalized shortfalls experienced by the chronically poor, or $|g(\tau)|$, divided by TN , which is the maximum value this sum can take.¹³

While K_1 is sensitive to magnitude of the income shortfalls of the chronically poor, the specific way the gaps are combined ensures that a given sized income decrement has the same effect on overall poverty whether the gap is large or small. One could argue that a loss in income would have a greater effect the larger the gap, in which case the *square* of the normalized gaps, rather than the gaps themselves, could be used. For example, suppose that the initial level of income is 4 and the poverty line is 5, so that the normalized gap is 0.20 and the squared (normalized) gap is 0.04. Decreasing the income by one unit will increase the squared gap to 0.16, an increase of 0.08. Now suppose that the initial level of income is 2, so that the normalized gap is 0.60 and the squared gap is 0.36. The unit decrement would raise the squared gap to 0.64, which represents a much larger increase of 0.28. Using squared gaps, rather than the gaps themselves, places greater weight on larger shortfalls.

Consider the matrix $s(\tau)$ whose nonzero entries are the squared normalized gaps of the chronically poor. The number of nonzero entries is $|h(\tau)|$ so that the *average squared gap* over these periods of poverty is given by $S(\tau) = |s(\tau)|/|h(\tau)|$. If this partial index is used instead of $G(\tau)$ to supplement the duration adjusted headcount ratio, the resulting chronic

¹³ Tn is the value of $|g(\tau)|$ that would arise in the extreme case where all incomes were 0.

poverty index would place greater weight on larger shortfalls. The resulting *duration adjusted FGT measure* $K_2 = K_0 S$ is a chronic poverty analog of the usual FGT index P_2 (just as K_0 and K_1 respectively correspond to P_0 and P_1 of the same class). K_2 has a straightforward expression as the mean of the entries of the matrix $s(\tau)$ of squared gaps: $K_2 = \mu(s(\tau)) = |s(\tau)|/(TN)$. It is the sum of the squared (normalized) gaps of the chronically poor, divided by the maximum value this sum can take.

Referring once again to the numerical example, the matrix of squared gaps is given by

$$s(\tau) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0.16 & 0.04 & 0 \\ 0 & 0.04 & 0.36 & 0 \\ 0 & 0.16 & 0.36 & 0 \end{bmatrix}$$

and hence $K_2 = \mu(s(\tau)) = (1.12)/16 = 0.07$. Now recall that the income of person 2 in period 3 is $y_2^3 = 4$, so that a unit decrement in income causes the squared normalized gap to rise from 0.04 to 0.16, and raising K_2 by about 0.008. In contrast, a unit decrement from $y_3^3 = 3$ raises the squared normalized gap from 0.16 to 0.36, and lifting K_2 by about 0.013. With K_2 , the impact of a unit decrement is larger for lower incomes than for higher incomes.

Analogous reasoning demonstrates that K_2 is sensitive to the distribution of income among the poor. Let i and j be two chronically poor persons with income vectors y_i and y_j . Suppose that their income vectors are replaced with $y_i' = \lambda y_i + (1-\lambda)y_j$ and $y_j' = (1-\lambda)y_i + \lambda y_j$, respectively, for some $\lambda \in (0, 1/2]$. This represents a uniform “smoothing” of the incomes of persons i and j , with the value $\lambda = 1/2$ yielding the limiting case where $y_i' = y_j' = (y_i + y_j)/2$ is a simple average of the two vectors. A transformation of this type is the multi-dimensional analog of a progressive transfer (among the poor) and it is easy to show that K_2 will not rise; indeed, if their associate normalized gap distributions g_i and g_j were not initially identical, K_2 would fall as result of the progressive transfer.¹⁴ In the numerical example, if the income vectors of persons 2 and 3 are replaced by the average vector (with $\lambda = 1/2$), then K_2 falls from 0.07 to about 0.54. In contrast, this smoothing of incomes affects neither the average duration of poverty, nor the average shortfall among the chronically poor, and hence K_0 and K_1 are entirely unaffected. The property requiring such a transformation to decrease chronic poverty is called the transfer axiom in the next section: K_2 satisfies this axiom while K_0 and K_1 violate it.

The general approach to constructing chronic poverty measures can be applied to obtain analogs of all of the indices in the FGT class. For any $\alpha \geq 0$ let $g^\alpha(\tau)$ be the matrix whose entries are the α powers of normalized gaps for the chronically poor (and zeros for those

¹⁴ See Kolm (1977) and Tsui (2002). The condition g_i and g_j rules out the case mentioned by Tsui (2002) where the two chronically poor persons are poor in the same periods and have the same incomes below the poverty line.

who are not chronically poor).¹⁵ The *duration adjusted* P_α measures are the general class of chronic poverty measures defined by $K_\alpha(y; z, \tau) = \mu(g^\alpha(\tau)) = |g^\alpha(\tau)|/(TN)$; in other words, K_α is the sum of the α power of the (normalized) gaps of the chronically poor, divided by the maximum value that this sum could take. It is an easy matter to define the associated class of *transient* poverty measures. Note that $K_\alpha(y; z, 0) = \mu(g^\alpha)$, so that when $\tau = 0$, the chronic poverty measure K_α takes into account every spell of poverty for all persons. The measure $K_\alpha(y; z, \tau) = \mu(g^\alpha(\tau))$ includes only the gaps of the chronically poor. Hence, $R_\alpha(y; z, \tau) = K_\alpha(y; z, 0) - K_\alpha(y; z, \tau)$ includes only the gaps of those who are not chronically poor, and hence is the associated measure of transient poverty. This definition will be used in our empirical application below. We now turn to a discussion of the properties satisfied by chronic poverty measures.

D. PROPERTIES FOR CHRONIC POVERTY INDICES

Our first basic axiom provides a formal way of ensuring that the income variable is comparable across persons. We say that $x \in D$ is obtained from $y \in D$ by a *permutation of incomes across people* if there exists an $N \times N$ permutation matrix¹⁶ P such that $x = yP$. A permutation of incomes changes the ordering of the vectors y_1, \dots, y_n in the distribution matrix, so that the income vector y_i previously received by person i in y is now received by a potentially different person j in x . The first property ensures that a measure of chronic poverty $K(x; z, \tau)$ is unaffected by such a transformation.

Anonymity If x is obtained from y by a permutation of incomes across persons, then $K(x; z, \tau) = K(y; z, \tau)$.

Under anonymity, the specific names (or index numbers) attached to the income vectors have no consequence for the measurement of chronic poverty.

The next basic axiom is designed to ensure that the measure makes coherent decisions across different sized populations. We say that x is obtained from y by a *replication of incomes across people* if there is some $M \geq 2$ such that $x \in D^N$ and $y \in D^{NM}$ where $x = (y, \dots, y)$. In other words, the matrix x is made up of M copies of each income vector in y , or more colloquially, each person in y has M “clones” in x . Once again, the requirement is that this form of transformation should leave chronic poverty unchanged.

Replication Invariance If x is obtained from y by a replication of incomes across people, then $K(x; z, \tau) = K(y; z, \tau)$.

Under this requirement, a replication of incomes across people leaves chronic poverty unchanged. In particular, it ensures that chronic poverty does not rise just because there are more people in one society than another; rather, chronic poverty is measured in per capita terms and, in this sense, is independent of population size.

The identification step has determined both the chronically poor group Z and those who are not chronically poor, namely, all i not in Z . One of the key properties for a chronic

¹⁵ For $\alpha = 0$, the entries of the matrix are more precisely defined as the limit of the entries of $g^\alpha(\tau)$ as α tends to 0.

¹⁶ A permutation matrix is a square matrix whose entries are “0” or “1” and each column and row sums to 1.

poverty measure is that it should not be sensitive to the income levels of those who are not identified as chronically poor. We say that x is obtained from y by a *simple increment to a nonpoor income* if there is some period t' , and a person i' who is not chronically poor in y , such that $x_{i'}^{t'} > y_{i'}^{t'}$ for $(i,t) = (i',t')$ and $x_i^t = y_i^t$ for all $(i,t) \neq (i',t')$. In other words, the two distributions x and y are only different for a single period's income for a person who is not chronically poor, and this income is larger in x than in y .

Focus If x is obtained from y by a simple increment to a nonpoor income, then

$$K(x;z,\tau) = K(y;z,\tau).$$

In other words, if a person is not chronically poor, then the specific incomes of that person should not be relevant for the measurement of chronic poverty. Note that this conclusion is intuitive in the case where the income in question is above the poverty line. But even when the income is *below* the poverty line, but the individual is not chronically poor, chronic poverty should not be altered by the increment.

In contrast, if an income of a chronically poor person falls during a spell of poverty, it is intuitive that the chronic poverty level should rise. We say that x is obtained from y by a *simple decrement to a poor income* if there is some period t' , and a person i' who is chronically poor in y , such that $x_{i'}^{t'} < y_{i'}^{t'} < z$ for $(i,t) = (i',t')$ and $x_i^t = y_i^t$ for all $(i,t) \neq (i',t')$. In other words, the two distributions x and y are only different for a single period's income for a person who is both chronically poor and in poverty during that period; and during this period the person's income is smaller in x than in y .

Monotonicity If x is obtained from y by a simple decrement to a poor income, then $K(x;z,\tau) > K(y;z,\tau)$.

This axiom takes the non-controversial position that chronic poverty should be sensitive to the actual income of a chronically poor person during a spell of poverty, and that chronic poverty should in fact rise as this income falls.

These are the four basic properties for chronic poverty measures.¹⁷ The first three are satisfied by all of the chronic poverty measures discussed here, as can be immediately seen by examining the matrices used to define each measure; monotonicity is just violated by the counting measures H and K_0 ; the definition $K_\alpha = \mu(g^\alpha(\tau))$ ensures that monotonicity is satisfied by K_α for $\alpha > 0$.

The next three properties concern the time dimension and are natural in the context of the spells approach to chronic poverty measurement (but may well be violated under other views of chronic poverty). We say that x is obtained from y by a *permutation of incomes across time* if there exists a $T \times T$ permutation matrix Π for which $x = \Pi y$. This type of

¹⁷ See also Tsui (2002), Kanbur and Mukherjee (2006), Dercon and Calvo (2006), and Foster and Santos (2006) for further discussions of these and other axioms. Kanbur and Mukherjee, in particular, have pointed out a problem that arises when replication invariance is satisfied in the presence of an epidemic like HIV-AIDS: a lower level of measured poverty may be arising not from poor persons being lifted out of poverty, but from their succumbing to the disease. Of course, measured poverty has indeed fallen in this case, but it is important for researchers to understand the underlying cause of this change.

transformation has the effect changing the timing of the income distributions y^1, \dots, y^T in the distribution matrix, so that the income distribution y^t previously received at date t in y is now received at date t' in x , with t' potentially being different from t . The next property ensures that a measure of chronic poverty $P(x; z, \tau)$ is unaffected by such a transformation.

Time Anonymity If x is obtained from y by a permutation of incomes across time, then $K(x; z, \tau) = K(y; z, \tau)$.

Under time anonymity, the ordering of the incomes does not affect the value of the chronic poverty measure. It is immediate that each of the chronic poverty measures we have presented satisfies this property; however, it is not entirely clear whether it is too severe a simplification. For example, it rules out the possibility that a “bunching” of periods in poverty may create greater harm than when the same periods in poverty are interspersed with non-poverty spells. It does not allow for the possibility that earlier incomes may have greater value, as typically expressed in the discounting of per-period incomes.¹⁸ Conversely, it could not accommodate a view of chronic poverty in which spells of poverty experienced in the more distant past have less salience than more recently experienced spells. It is not entirely clear whether and how the time-ordering of incomes should impact the aggregation (or identification) of chronic poverty; this property takes the extreme position that the measure should ignore the time-ordering entirely.¹⁹

The focus axiom for static poverty measures is based on an implicit assumption that (apart from within the household unit) incomes from one person cannot be transferred to the incomes of another person. This allows the set of the poor to be well defined and ensures the measure will ignore the incomes of persons outside this set. In an analogous way, the next form of focus axiom relies on an implicit assumption that incomes cannot be transferred across periods (either physically or conceptually, as envisioned by compensation principles). Hence, an increase in income during a period when a chronically poor person is not in poverty will not serve to lower chronic poverty at all. We say that x is obtained from y by a *simple increment to a nonpoor income of a chronically poor person* if there is some period t' , and a person i' who is chronically poor in y , such that $x_i^{t'} > y_i^{t'} \geq z$ for $(i, t) = (i', t')$ and $x_i^t = y_i^t$ for all $(i, t) \neq (i', t')$. In other words, the two distributions x and y only differ in only a single income – when a chronically poor person is not in a spell of poverty – and this particular income is larger in x than in y .

¹⁸ See Rodgers and Rodgers (1993) or Dercon and Calvo (2006) for examples of chronic poverty measures using discounting.

¹⁹ We might also define the analog of the replication invariance axiom in the time dimension by stacking a given matrix y vertically m times (resulting in a *replication of incomes across time*) and requiring the measure to be independent of such transformations. Each of the measures defined above, when appropriately extended to the domain of income matrices having an arbitrary number of periods, would satisfy this property. For simplicity of presentation, we have fixed the number of time periods at T for this paper and will not be exploring the advisability of such a property.

Time Focus If x is obtained from y by a simple increment to a nonpoor income of a chronically poor person, then $K(x; z, \tau) = K(y; z, \tau)$.

In other words, if a person is chronically poor but not currently in a spell of poverty, the measure should ignore the current level of income. Under certain conceptions of chronic poverty it may make sense to take into account all the income levels of a chronically poor person. However, in the spells approach to measuring chronic poverty, each period's income is not directly aggregated with the next, and hence the time focus axiom is a natural requirement.

The above property makes a sharp distinction between periods when a chronically poor person is in poverty, and periods when the person is not. Given this, it is natural to regard the time spent in poverty to be an important aspect of chronic poverty. The next property is a general requirement that the measure be sensitive to the time a chronically poor person spends in poverty. We say that x is obtained from y by a *duration enhancing decrement to a chronically poor person* if there is some period t' , and a person i' who is chronically poor in y , such that $x_i^{t'} < z \leq y_i^{t'}$ for $(i, t) = (i', t')$ and $x_i^t = y_i^t$ for all $(i, t) \neq (i', t')$. In other words, the two distributions x and y are only different for a single period's income for a person who is chronically poor, and for that period the person is not in poverty in y , but falls into poverty in x .

Time Monotonicity If x is obtained from y by a duration enhancing decrement to a chronically poor person, then $K(x; z, \tau) > K(y; z, \tau)$.

During a period in which a chronically poor person happens to be having a spell *outside of poverty*, if the income level falls below the poverty line (thus raising the number of duration of poverty experienced by this person), then poverty should rise.

These are the three axioms concerned with the time element of chronic poverty. By construction all of the measures we have considered satisfy both the time anonymity and time focus axioms. While H violates time monotonicity, each K_α measure satisfies this axiom since a duration enhancing decrement to a chronically poor person adds a positive entry to $g^\alpha(\tau)$ and hence $K_\alpha = \mu(g^\alpha(\tau))$ must rise.

The next property provides a chronic poverty analog of the transfer axiom of Sen (1976), which require the poverty measure to take into account inequality among the poor. In that environment, each poor person has a single income and consequently a “smoothing” of incomes may be viewed as a progressive transfer; in the present case, chronically poor people can have several periods with incomes below z , and a natural extension is given by Kolm's (1977) multidimensional smoothing transformation which uses an $N \times N$ bistochastic matrix²⁰ B to obtain a distribution $x = yB$ in which the incomes in each period are “averaged” in the same way. Two restrictions on x and y must be specified in the present context. First, the income vectors of all persons who are not chronically poor are to be left unchanged. Second, x and y cannot have $g(\tau)$ matrices that are permutations of one another (which would happen if, say, the transformation only covered chronically poor persons with identical distributions). Consider the following definition that rules

²⁰ A bistochastic matrix is a square matrix whose entries are nonnegative and each column and row sums to 1.

out these exceptional cases: We say that x is obtained from y by a *smoothing of chronically poor incomes* if there is an $N \times N$ doubly stochastic matrix B such that $x = yB$, where (i) $x_i = y_i$ for all i not in Z , and (ii) the $g(\tau)$ matrix associated with x is not a permutation of the $g(\tau)$ matrix for y .

Transfer If x is obtained from y by a smoothing of chronically poor incomes, then $K(x; z, \tau) < K(y; z, \tau)$.

Thus, a smoothing of chronically poor incomes should lower chronic poverty.²¹

The final two properties relate chronic poverty levels within population subgroups to chronic poverty overall.

Subgroup Consistency Suppose that $P(x; z, \tau) < P(y; z, \tau)$ and $P(x'; z, \tau) = P(y'; z, \tau)$ for distributions x, y, x' , and y' satisfying $N(x) = N(x')$ and $N(y) = N(y')$. Then $P(x, x'; z, \tau) < P(y, y'; z, \tau)$.

This property requires that if poverty falls in given subgroup and stays fixed in the remaining subgroup (where the population sizes of the two subgroups are unchanged), then overall poverty must fall. By repeated application, its scope expands to cover the case where both local poverty levels fall, or where there are many subgroups. The motivation of this property is to ensure coherence between local and global evaluations of poverty, which is a basic requirement for targeted poverty alleviation policies.

A second property linking subgroup to overall poverty not only requires coherence, it specifies the precise formula by which this occurs.

Decomposability For any distributions x and y we have

$$K(x, y; z, \tau) = \frac{N(x)}{N(x, y)} K(x; z, \tau) + \frac{N(y)}{N(x, y)} K(y; z, \tau).$$

According to this property, when a distribution is broken down into two subpopulations, the overall chronic poverty level can be expressed as a weighted average of subgroup chronic poverty levels, with the weights being the respective subgroup population shares.

These are the three “higher order” properties for chronic poverty measures that are somewhat more specialized in nature. All of the measures we have discussed satisfy population decomposability and hence subgroup consistency; this is expected given that the static versions also satisfy these properties. Consequently, H and K_α are well suited for the analysis of chronic poverty by population subgroup. It is easy to show that H and K_α for $\alpha \leq 1$ violate the transfer axiom; to see that K_α for $\alpha > 1$ satisfies the axiom, we note that a smoothing of chronically poor incomes leads to a gap matrix $g(\tau)$ that can be obtained from the original gap matrix by a bistochastic matrix and by a series of simple decrements. Since for $\alpha > 1$ the measure $K_\alpha = \mu(g^\alpha(\tau))$ is a convex, increasing function

²¹ Notice that when chronically poor persons differ their timing of poverty spells, the smoothing process may entail a transfer from an income above z to one that is below. In the presence of the time focus axiom and monotonicity, we already know that chronic poverty would be lowered as the higher income fell to z , so there is some overlap between them and the transfer axiom. One could instead restrict the axiom to cases where the affected chronically poor persons were poor in the same periods, yielding a *weak transfer axiom* (as is often defined in the static environment).

of the entries of $g(\tau)$, it follows that K_α must fall as a result of this transformation, hence K_α satisfies the transfer axiom for $\alpha > 1$. In particular, the chronic analog of the squared gap measure K_2 satisfies all of the properties defined above.

IV. AN EMPIRICAL ILLUSTRATION

In this section, the chronic poverty measures are applied to panel data from Argentina's Instituto Nacional de Estadísticas y Censos (INDEC), covering 2409 households during the four waves of October 2001, May 2002, October 2002, and May 2003. The income variable used is the equivalent household income, calculated by dividing total household income in each period by the number of equivalent adults (using the equivalent adult scale provided by INDEC).²² The Instituto also provides a separate poverty line for each region and each period in order to capture spatial variations in the cost of living as well as inflation over time. Normalized gaps for a household are found by subtracting the equivalent income from the appropriate poverty line and dividing this difference by the same line. Since there are $T = 4$ periods, there are four relevant duration cutoffs τ , namely, 0.25, 0.50, 0.75, 1.00. Table 1 provides the resulting levels of chronic poverty for each of the measures H , K_0 , K_1 , and K_2 , at each of the four duration cutoffs. Notice that for each index, the measured level of chronic poverty rises as the duration requirement τ falls. For $\tau = 1$, the headcount and duration adjusted headcount have the same value since the average duration in this special case is precisely 1. In general, for fixed τ the value of H is higher than that of K_0 , which in turn is higher than K_1 , and so forth, reflecting the mathematical properties of the measures. However, it is important to remember that actual poverty comparisons should only be made using the same poverty measure and with the same duration cutoff.

The first column provides the chronic poverty level when the duration in poverty is at least 0.25, and this is clearly the same as the level when the minimum duration is 0. This latter level $K_\alpha(y; z, 0)$ is precisely the static poverty level that would arise if the data from all periods were merged into a single distribution and evaluated by P_α , or equivalently, $\Sigma_t P_\alpha(y^t; z) / \tau$ the average static poverty level across all periods. Now given a minimum duration of $\tau = 0.75$, the difference between this overall poverty level $K_\alpha(y; z, 0)$ and the chronic poverty level $K_\alpha(y; z, \tau)$ is the level of transient poverty, $R_\alpha(y; z, \tau)$ as defined above. The final two columns provide the percentage of overall poverty due to chronic poverty and transient poverty respectively for $\tau = 0.75$ given each measure. It is interesting to note that as the poverty measure changes from H to K_0 to K_1 to K_2 , the share of chronic poverty becomes larger while the share of transient poverty falls. For K_2 the chronic poverty share is 90.6%, suggesting that the transient poverty spells generally do not involve large shortfalls.

²² In periods where an incomes is missing or zero for a given household, the methodology of Little and Su (1989) is followed, which calculates an imputed income using the household's incomes at other dates and other household incomes at the same date. In the few cases where the household reports zero income in all four periods, an income equal to the lowest social welfare transfer was assumed.

Table 2 provides a profile of chronic poverty in each of the six regions of Argentina, namely, Greater Buenos Aires (GBA), the Northeast (NE), the Northwest (NW), the Center of the country (C), the Midwest (MW), and the South (S). In this example, chronically poor means poor in all four periods; hence the duration cutoff is $\tau = 1$. The first column in the table gives the population shares of the six regions. The second, fourth and sixth columns provide the regional chronic poverty levels for K_0 , K_1 , and K_2 , respectively; while the third, fifth and seventh columns gives the percentage contribution of each region to total poverty (or the population share times the regional poverty level over the total poverty). Notice that the Northeast has only 15% of the population, yet accounts for at least 26% of the chronic poverty in Argentina; in contrast, the South has 12% of the total population, but only contributes 4% or less to total chronic poverty.

V. CONCLUSIONS

This paper has presented a new family K_α of chronic poverty measures based on the P_α measures of Foster, Greer and Thorbecke (1984). Each measure has an intuitive interpretation as a “duration adjusted P_α measure” and can be readily calculated as the mean of a particular matrix. All the measures satisfy a series of general axioms for chronic poverty measures, and an additional set of axioms that are applicable to measures based on a spells approach to chronic poverty. Within the class, there is a range of measures satisfying a monotonicity axiom, and a smaller subset satisfying a transfer axiom. In particular, K_2 , the duration-adjusted analog of the usual P_2 measure, satisfies all of the properties given in this paper. The usefulness of the class of chronic poverty measures has been illustrated using panel data from Argentina.

One powerful criticism of the general framework for evaluation used here is that, being limited to a single income variable, it cannot utilize data on other dimensions of relevance to poverty and wellbeing. The approach is indeed restrictive, and when panel data containing information on other capabilities become more widely available, multidimensional measures of chronic poverty will need to be developed. Of course, there remain serious difficulties in formulating effective multidimensional poverty measures, even in the static context. The identification step for multidimensional poverty depends crucially on assumptions about substitutability across different dimensions; and the aggregation step has seemingly endless possibilities for bringing the various dimensions of deprivation into a single coherent index. Even the process of making the dimensions commensurate requires a leap of faith. The problem of measuring multidimensional chronic poverty is not likely to be solved anytime soon.

The measurement approach presented in this paper is based on a very simple treatment of time in poverty: an earlier period in poverty is given the same weight as a later period in poverty. Indeed, such a position could be justified using the traditional “no discounting” argument of Ramsey (1928). However, other plausible alternatives exist. For example, one might argue that greater weight should be placed on income received in *earlier* periods, as is typically done by discounting (see Rodgers and Rodgers, 1993, or Dercon and Calvo, 2006). This could be justified *ex ante* where there may be greater perceived or actual value from receiving income earlier. However, when viewed from an *ex post*

perspective, there may also be good reason to view earlier incomes as having less weight (see, for example, the discussion in Ray and Wang, 2001). An additional problem with discounting (one way or the other) is deciding *what* to discount: the income received by the chronically poor person, the utility level, or some function of the gap. And if the discounting procedure delivers a different ordering depending on the base period for evaluation, this could certainly make the procedure less attractive.

There are other aspects of the above approach that might be altered. Some have argued that continuous spells of poverty should have larger weight in the evaluation of chronic poverty. As noted by x, there are significant data problems from trying to implement this using panel data. Another perspective might consider placing greater weight on observations from a chronically poor person who is in poverty longer. This could indeed be done either via a new functional form or through dominance orderings over time, but this will have to await future work.

REFERENCES

- Atkinson, A. B. (1987) "On the Measurement of Poverty", *Econometrica* 55, 749-64.
- Bane, M. and D. Ellwood (1986), "Slipping into and out of Poverty: The Dynamics of Spells", *Journal of Human Resources*, 21, 1-23.
- Brooks-Gunn, J. and G. J. Duncan (1997), "The Effects of Poverty on Children", *Children and Poverty*, 7, 55-70.
- Chakravarty, S. (1983), "A New Index of Poverty", *Mathematical Social Sciences* 6, 307-313.
- Chronic Poverty Research Centre (2005), *The Chronic Poverty Report 2004-05*, University of Manchester.
- Clark, D. and D. Hulme (2005), "Towards A Unified Framework for Understanding the Depth Breadth and Duration of Poverty", Mimeo, Manchester.
- Clark C. R., R. Hemming and D. Ulph (1981), "On Indices for the Measurement of Poverty", *Economic Journal* 91, 515-26
- Dercon, S. and S. Calvo, "Chronic Poverty and All That: The Measurement of Poverty Over Time", mimeo, University of Oxford, 2006.
- Foster, J. E. (2005), "Poverty indices" in *Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke*, Alain de Janvery and Ravi Kanbur, (eds), forthcoming.
- Foster, J. E., J. Greer and E. Throbecke (1984), "A Class of Decomposable Poverty Indices", *Econometrica*, 52, 761-6.
- Foster, J. E. and M. A. Santos (2006), "Measuring Chronic Poverty", Mimeo, Vanderbilt.
- Foster, J. E. and A. Sen (1997) "On Economic Inequality: After a Quarter Century" Annex to the enlarged edition of *On Economic Inequality*, by Amartya Sen, Clarendon Press.
- Foster, J. E. and A. F. Shorrocks (1988a), "Poverty Orderings and Welfare Dominance" *Social Choice and Welfare*, 1988, 179-98.
- Foster, J. E. and A. F. Shorrocks (1988b), "Subgroup Consistent Poverty Indices" *Econometrica*, 59, 687-710.
- Foster, J. E. and M. Székely "Poverty Lines over Space and Time" mimeo, Vanderbilt University, October 2006.
- Gaiha, R. and A. B. Deolaiker (1993), "Persistent, Expected and Innate Poverty: Estimates for Semi Arid Rural South India". *Cambridge Journal of Economics*, 17, 409-421.
- Hulme, D., and Shepherd, A., 2003, "Conceptualising Chronic Poverty", *World Development*, 1(3), 403-425.
- Hulme, D., K. Moore and A. Shepherd (2001), "Chronic Poverty: Meanings and Analytical Frameworks", *CPRC working paper 2*, Manchester: IDPM, University of Manchester.
- Jalan, J. and M. Ravallion (1998), "Determinants of Transient and Chronic Poverty, Evidence from rural China", *Policy Research Working Paper 1936*, The World Bank.
- Kanbur, R. and D. Mukherjee (2006), "Premature Mortality and Poverty Measurement", Mimeo, Cornell University.
- Kolm, S.-Ch. (1977), "Multidimensional Egalitarianisms", *Quarterly Journal of Economics*, 91, 1-13.

- Lichter, D. T., M. J. Shanahan, and E. L. Gardner (1999), "Becoming a Good Citizen? The Long-Term Consequences of Poverty and Family Instability During Childhood", Mimeo, Russell Sage.
- Little, R. J. A. and H.-L. Su (1989), "Item Non-Response in Panel Surveys" in D. Kasprzyk, G. Duncan and M. P. Singh (Eds.), *Panel Surveys*, New York: John Wiley.
- McKay, A. and D. Lawson (2002), "Chronic Poverty: A Review of Current Quantitative Evidence", CPRC Working Paper No. 15, Manchester: University of Manchester.
- McKay, A. and D. Lawson (2003), "Assessing the Extent and Nature of Chronic Poverty in Low Income Countries: Issues and Evidence", *World Development*, 31, 425-439.
- McDonough, P. and P. Berglund (2003), "Histories of Poverty and Self-Rated Health Trajectories", *Journal of Health and Social Behavior*, 44, 198-214.
- Ray, D. and R. Wang (2001), "On Some Implications of Backward Discounting" mimeo, New York University.
- Rodgers, J. R. and J. L. Rodgers (1993), "Chronic Poverty in the United States," *Journal of Human Resources*, 18, 25-54.
- Ramsey, F. (1928), "A Mathematical Theory of Savings," *Economic Journal* 38, 543-59.
- Ravallion, M. (1994), *Poverty Comparisons*, Chur, Switzerland: Harwood.
- Sen, A., (1976), "Poverty: An Ordinal Approach to Measurement", *Econometrica*, 44, 219-231.
- Shorrocks, A. F. (1992), "Spell Incidence, Spell Duration and the Measurement of Unemployment", *Working Papers of the ESRC Research Centre on Micro-social Change*, Paper N° 12, University of Essex.
- Stevens, A.H. (1994), "The Dynamics of Poverty Spells: Updating Bane & Ellwood", *American Economic Review - Papers and Proceedings*, 84, 34-37.
- Tsui, K.-y. (2002), "Multidimensional Poverty Indices", *Social Choice and Welfare*, 19, 69-93.
- Yaqub, Shafin. (2000), "Poverty Dynamics in Developing Countries", IDS Development Bibliography, University of Sussex, April.
- Yaqub, S. (2001), "At What Age Does Poverty Damage Most?" Mimeo, Sussex.
- Yaqub, S. (2003), "Chronic Poverty: Scrutinizing Estimates, Patterns, Correlates and Explanations", CPRC Working Paper 21, University of Manchester.

TABLE 1: CHRONIC POVERTY IN ARGENTINA: ESTIMATED LEVELS FOR VARIOUS MEASURES AND DURATIONS

	MEASURE			
	H	K ₀	K ₁	K ₂
$\tau = 0.25$	0.47	0.34	0.15	0.085
$\tau = 0.50$	0.38	0.31	0.14	0.082
$\tau = 0.75$	0.30	0.27	0.13	0.077
$\tau = 1.00$	0.20	0.20	0.10	0.064
%Chronic ^a	63.8%	79.4%	86.7%	90.6%
%Transient	36.2%	20.6%	13.3%	9.4%

^aDuration cutoff is $\tau = 0.75$. Percentage Chronic is found by dividing third row by first row; Percentage Transient is 1 minus Percentage Chronic.

TABLE 2: CHRONIC POVERTY PROFILE FOR ARGENTINA

	MEASURE^a						
	Pop Share	K₀	Percent Contrib	K₁	Percent Contrib	K₂	Percent Contrib
GBA	12%	0.13	8%	0.07	8%	0.042	8%
NE	15%	0.37	26%	0.19	27%	0.122	28%
NW	25%	0.26	29%	0.13	28%	0.077	27%
C	28%	0.17	24%	0.09	25%	0.058	25%
MW	11%	0.18	10%	0.09	9%	0.050	9%
S	12%	0.06	4%	0.03	3%	0.016	3%
Total	100%	0.20		0.10		0.064	

^a Duration cutoff is $\tau = 1.00$. Percentages may not sum to 100% due to rounding errors.