Institute for International Economic Policy Working Paper Series
Elliott School of International Affairs
The George Washington University

The Foster-Greer-Thorbecke (FGT) Poverty Measures: Twenty-Five Years Later

IIEP-WP-2010-14

James E. Foster
George Washington University

Joel Greer

Erik Thorbecke
Cornell University

April 2010

Institute for International Economic Policy
1957 E St. NW, Suite 502
Voice: (202) 994-5320
Fax: (202) 994-5477
Email: **iiep@gwu.edu**
Web: **www.gwu.edu/~iiep**

The Foster-Greer-Thorbecke (FGT) Poverty Measures: Twenty-Five Years Later

by

James Foster*

Joel Greer†

Erik Thorbecke§

*The Elliott School of International Affairs and the Department of Economics, The George Washington University, Washington, DC 20052, USA, and Oxford Poverty & Human Development Initiative, Oxford Dept of International Development, University of Oxford, QEH, 3 Mansfield Road, Oxford OX1 3TB, UK, e-mail: fosterje@gwu.edu

†Silver Spring, Maryland

§Department of Economics and Division of Nutritional Sciences, B-16 MVR, Cornell University, Ithaca, New York 14853, e-mail: et17@cornell.edu. (for correspondence and/or offprints).

Date of Receipt of Original Manuscript:

Date of Acceptance of Revised Manuscript:

Abstract Twenty-five years ago, the FGT class of decomposable poverty measures was introduced in Foster, Greer, and Thorbecke [54]. The present study provides a retrospective view of the FGT paper and the subsequent literature, as well as a brief discussion of future directions. We identify three categories of contributions: to measurement, to axiomatics, and to application. A representative subset of the literature generated by the FGT methodology is discussed and grouped according to this taxonomy. We show how the FGT paper has played a central role in several thriving literatures and has contributed to the design, implementation, and evaluation of prominent development programs: the breadth of its impact is evidenced by the many topics beyond poverty to which its methodology has been applied. We conclude with a selection of prospective research topics.

Keywords Axioms, decomposability, FGT measures, income distribution, poverty, stochastic dominance, subgroup consistency

Acknowledgments The authors thank Ravi Kanbur for suggesting that we write this paper; Sabina Alkire, Chris Bennett, and Stephen Smith for helpful comments and encouragement; and Oswaldo Molina-Campodonico, Mark Eisinger, Suman Seth, and Shabana Singh for research assistance. We owe a debt of gratitude to Ravi Kanbur and Martin Ravallion, whose research and recommendations have introduced the FGT measures to a much broader audience. Foster is grateful to the Elliott School's Institute for International Economic Policy for research funding.

1 Introduction

Twenty-five years ago, we introduced a class of poverty measures having the formula $P_\alpha = \frac{1}{n}\Sigma_{i=1}^{q}(\frac{z-y_i}{z})^\alpha$ where z is the poverty line, y_i is the i^{th} lowest income (or other standard of living indicator), n is the total population, q is the number of persons who are poor, and $\alpha \geq 0$ is a "poverty aversion" parameter. We recognized at the time that the class has certain advantages. Its simple structure—based on powers of normalized shortfalls—facilitates communication with policymakers. Its axiomatic properties are sound and include the helpful properties of additive decomposability and subgroup consistency, which allow poverty to be evaluated across population subgroups in a coherent way. Initial responses from colleagues suggested that our new tools might well contribute to the literature on poverty.

However, as time has unfolded, we have been surprised by the impact of the work on theory, application, and policy in the domain of poverty evaluation, and other unrelated domains. For instance, subsequent research has concluded that the indices are closely linked to stochastic dominance and provide a unifying structure linking poverty, inequality, and wellbeing. The measures have become the standard for international evaluations of poverty, are reported regularly by the World Bank's PovcalNet, by a host of UN agencies, and by individual countries. They are naturally suited for targeting exercises and other policy implementations. They have a central role in the growing statistical literature on stochastic dominance and multidimensional dominance tests. And they have been adapted to measure a host of other phenomena, such as the "graying" of a population, corruption, obesity, the "rich," affordability of low-income housing, food insecurity, and the research productivity of economics departments. One crude indicator of the impact of the paper is its citation profile over the intervening 25 years. Data from the *Social Sciences Citation Index* reveals a steadily increasing number of annual citations, rising at an average rate of about 1.5 per year. The cumulative count is now 506, with 126 citations in the last three years alone.[1] More expansive databases, such as Google Scholar, count citations in the thousands. Other authors apply FGT measures without formal attribution, suggesting that they have moved into the realm of the generic.

Now that a quarter century has passed since the publication of our paper, we thought that it could be an opportune time to reflect upon some of its main themes. The resulting essay is largely retrospective—constituting a survey of subsequent developments—but also includes a shorter prospective exploration. The paper has been a "labor of love" for the coauthors, and we hope that those who are familiar with our original effort will be well entertained, while those unfamiliar will be motivated to explore the literature even further.

We begin in Section 2 with a brief description of the context and origin of the paper, and then turn to an assessment of its contributions—to measurement, axiomatics, and empirical work. Section 3 reviews the subsequent work in each of these areas. Section 4 is devoted to a brief discussion of the road ahead, while Section 5 concludes.

2 Context, Origin, and Contributions

By the end of the 1970s, the toolkit for evaluating inequality had been substantially enriched by the work of Atkinson, Kolm, Sen, and many other contributors.[2] The literature was motivated by very practical questions, but was strongly influenced by the theoretical methods and substantive approaches of social choice theory, welfare economics, and risk analysis. A compelling axiomatic framework had been created for inequality measurement, and it was

---

[1] Citation data as of November, 2009. Another indicator is the presence of the measures in standard textbooks in development and public economics (Ray [107], Fields [51], Todaro and Smith[133], and Myles[100]) and in other specialized books and collections (Ravallion [103], Deaton [40], Subramanian [124], and Barrett [12]).

[2] See Atkinson [7]; Kolm [88], [89], [90]; and Sen [117], the surveys of Foster and Sen [61] and Cowell [36], and the retrospective of Kanbur [85].

influencing the way these concepts were being measured in practice. In turn, practical considerations were feeding back to the construction of new indices. Bourguignon [16], Shorrocks [120], and Cowell [33] converted Theil's [127] forms of additive decomposition into axioms, which were used to derive and characterize the generalized entropy class of inequality measures, including Theil's two measures.

In contrast to the active research on the measurement of inequality, the discussion of poverty measurement in the early 1970s largely concerned the selection of a "right" poverty line to identify the poor. Sen [118] brought into clear focus the importance of a second step in evaluating poverty—the aggregation step. He constructed an axiomatic framework for poverty measures, including two "dominance" axioms: a monotonicity axiom (that requires poverty to rise when a poor income falls), and a transfer axiom (that requires poverty to rise when a poor person transfers income to a richer poor person). The latter requirement—that the measure should reflect the distribution among the poor—was adapted from an analogous axiom in inequality analysis. Sen presented a new poverty measure and described a useful three-step procedure for deriving it. One of the steps assumes poverty to be a normalized weighted sum of shortfalls; a second step selects weights based on the rank order of poor incomes (an "ordinal" approach); a third indirectly sets the normalization factor. He justified the weights using an argument from the literature on relative deprivation. The resulting "Sen measure" can be expressed as

$$S(x;z) = H(I + (1-I)G_p)$$

where x is the income distribution, z is the poverty line, H is the headcount ratio or frequency of the poor, I is the income gap ratio or the average normalized shortfall among the poor, and G_p is the Gini coefficient among the poor.[3] The Sen measure takes into account the depth and distribution of the poor, in addition to the frequency. It reduces to HI (which has come to be

[3] See also Shorrocks [122] and Foster and Sen [61] for a revised Sen measure along the lines of Thon [128].

known as the poverty gap) in case there is no inequality among the poor; it rises towards H as G_p (or I) tends to 1.

Initially, much of the impact of Sen's paper was seen in theoretical discussions; empirical applications using the measure were less common. One exception was a thoughtful piece by Anand [3] on poverty in Malaysia, which employed the Sen measure and identified practical issues with its use. First, he noted that the poverty numbers rendered by the measure might not convey information in the way that, say, H does; instead, levels derive their meaning through comparison with other levels.[4] Second, he observed that the measure is not decomposable across subgroups, thus limiting its usefulness in analyzing regional data. When he constructed what he called a "profile of poverty" (following Orshansky [101]), he was forced to revert to the headcount ratio. Anand's [3] paper and subsequent 1983 book [4] exemplified the latent demand for a measure consistent with the Sen axioms, and yet more broadly applicable than the Sen measure.

This, then, is the context in which Foster, Greer, Thorbecke [54] was written. The coauthors, who were at Cornell University at the time, were made up of a theorist (Foster) and two empirically inclined development economists (Greer and Thorbecke). Thorbecke was concerned with evaluating food poverty in Kenya as part of a major project sponsored by its Ministry of Finance. Greer was a Ph.D. student who was part of the Kenyan project and whose dissertation under the direction of Thorbecke would address both conceptual and empirical issues of poverty measurement.[5] Foster was a Ph.D. student who had worked in axiomatic social choice theory and whose dissertation under the direction of Mukul Majumdar would analyze Theil's decomposable inequality measure. A visit by Amartya Sen to Cornell in October 1979 turned out

[4] A similar observation can be made for members of P_α; this general issue is discussed at greater length in Section 4.2 below.

[5] An expanded version of the FGT paper appeared as a chapter in Greer's dissertation.

to be an important catalyst for the FGT project. During the visit Thorbecke expressed doubts about Sen's "rank weighting " approach and proposed an alternative "shortfall weighting" method that became the basis for the squared gap measure P_2 featured in FGT paper. By early 1981, Cornell Department of Economics Working Paper No. 243 was circulated, which included several elaborations (such as the Pen's Parade diagram given below) that were later excised in the published version. The key concept of "subgroup monotonicity" (now called "subgroup consistency") was absent from earlier versions, but appeared in the final draft in response to a referee's call for a conceptual justification for additive decomposability. The paper was accepted in 1983 and appeared in 1984.[6]

The FGT paper sought to contribute to the literature on poverty in several ways: (i) it introduced a new class of poverty measures that is understandable, theoretically sound, and applicable; (ii) it helped justify the measures using new and practical axioms; and (iii) it provided a concrete illustration of the new technology. We now describe these contributions in greater detail.

---

[6] Its acceptance was perhaps surprising given the initial responses of the referees and editor. One referee expressed doubt "whether this contribution is sufficient to warrant publication" and the second was not particularly enthusiastic: "I tend to agree with the authors that their poverty measure…has obvious advantages over some of the other proposed measures." The editor remarked, "I am not anxious to publish 'another poverty index'" and "I would put the probability at around one half of acceptance."

2.1 Measurement

The FGT class is based on the *normalized gap* $g_i = (z-y_i)/z$ of a poor person i, which is the income shortfall expressed as a share of the poverty line.[7] Viewing g_i^α as the measure of *individual poverty* for a poor person, and 0 as the respective measure for non-poor persons, P_α is the average poverty in the given population. The case $\alpha = 0$ yields a distribution of individual poverty levels in which each poor person has poverty level 1; the average across the entire population is simply the headcount ratio P_0 or H. The case $\alpha = 1$ uses the normalized gap g_i as a poor person's poverty level, thereby differentiating among the poor; the average becomes the poverty gap measure P_1 or HI. The case $\alpha = 2$ squares the normalized gap and thus weights the gaps by the gaps; this yields the squared gap measure P_2. As α tends to infinity, the condition of the poorest poor is all that matters.

The parameter α has an interpretation as an indicator of "poverty aversion" in that a person whose normalized gap is twice as large has 2^α times the level of individual poverty. Alternatively, α is the elasticity of individual poverty with respect to the normalized gap, so that a 1% increase in the gap of a poor person leads to an α% increase in the individual's poverty level.[8] The parametric class of measures gave analysts and policymakers an instrument to evaluate poverty under different magnifying glasses with varying sensitivity to distributional issues.

The FGT paper emphasized the squared gap measure P_2, noting its simplicity and the fact that many arguments used in support of Sen's measure also apply to P_2. Sen had used a general

---

[7] Donaldson and Weymark [44] distinguish between a "weak" definition of the poor, which takes the poverty line as the minimum *non-poor* income, and a "strong" definition in which z is the maximum *poor* income. The original FGT paper followed Sen [118] and used the strong definition; here we use the weak definition.

[8] The interpretation of the parameter clearly drew its inspiration from the Atkinson [7] class of inequality measures and its "inequality aversion" parameter.

additive form for poverty measures in which poverty is a normalization factor times the weighted sum of the normalized gaps of the poor. He used rank orders as weights—so that the poorest person in a population of q poor persons is assigned a weight of q, the next has a weight of q-1, and so forth until the least poor person is assigned a weight of 1. Although this approach has been used to great effect in social choice theory (especially in the Borda count method), in the present context it could be viewed as being rather unresponsive to one's own condition and perhaps too responsive to the precise conditions of the other poor persons. By contrast, P_2 simply weights an individual's normalized gap by the normalized gap.

Sen argued that poverty should have an absolute deprivation component (as represented by the normalized gap) and a relative deprivation component (given by the weight) and then, invoking an example from the relative deprivation literature, posited that the latter was naturally represented by the rank order of a person in the group of the poor. One could argue, however, that this is not the only conception of relative deprivation. For example, an alternative notion of the *magnitude* of relative deprivation, which is "the extent of the difference between the desired situation and that of the person desiring it" [112], suggests the use of the normalized gap itself. Hence, it could be argued that P_2 is also composed of absolute and relative components.

The Sen measure could be expressed in terms of H, I, and G_p—where each component provides relevant information on the frequency, depth, and distribution of poverty.[9] The squared gap measure can likewise be expressed as $P_2 = H[I^2 + (1-I)^2C_p^2]$, where C_p^2 is the squared coefficient of variation (a decomposable inequality measure from the generalized entropy class) among the poor. Figure 1 from the 1981 working paper [53] provides a graphical representation of the link between C_p^2 and P_2 via this formula. A distribution is depicted in the style of Pen's Parade, but with ordered incomes given in poverty units and the total population normalized to 1. The quantities H, I, and 1-I are as noted, while C_p^2 is just the variance of the poverty-unit

[9] Foster and Sen [61] use the term "partial indices" to describe the components, since none is a fully-fledged aggregate measure of poverty satisfying the basic axioms for poverty measurement.

incomes of the poor about the mean. Suppose, instead, that all poor households had this mean level. Then each normalized shortfall g_i would be identical to I, and P_2 would be $\Sigma_{i=1}^{q} I^2 / n = HI^2$, as given by the formula when $C_p^2 = 0$. However, since the poor incomes are distributed unequally, P_2 is larger by $H(1-I)^2C_p^2$. When H and I are held constant, P_2 varies with C_p^2 according to this expression. This intuitive formula and the other justifications paralleling Sen helped with the interpretation of the new measure and gave it a certain credence. However, the *properties* satisfied by P_2 and the rest of the FGT class set it apart from its predecessor and gave it broader applicability.[10]

2.2 Axioms

Axioms for poverty measures can be usefully grouped under three general headings: invariance, dominance, and subgroup axioms [58]. The invariance category includes symmetry, replication invariance, scale invariance, focus, and continuity.[11] All the P_α measures satisfy the invariance axioms, with the exception of P_0 which has a discontinuity at the poverty line.

---

[10] Ideas do not appear in isolation, and it is not surprising that the FGT paper overlapped with others that appeared (or resurfaced) in the early 1980s. Kundu [91] did not mention decomposability or the P_α class, but defined P_2 and derived several of its properties. Clark, Hemming and Ulph [30] proposed two classes: one that can be converted to an additive decomposable class which includes Chakravarty's [23] family; a second equivalent to $H(P_\alpha)^{1/\alpha}$ for $\alpha > 1$. Foster [55] noted that Watts [135] had proposed a distribution sensitive poverty measure many years before Foster, Greer, and Thorbecke [53], or even Sen [118].

[11] Each invariance axiom states that poverty should be unaffected by some change in incomes: for symmetry, it is a permutation or switching of incomes; for replication invariance, it is a replication or cloning of the entire population; for scale invariance, it is a scaling up or down of all incomes and the poverty line; for the focus axiom, it is a change in a non-poor income which

The dominance category includes various versions of monotonicity, the transfer axiom, and transfer sensitivity.[12] Higher α reflects greater aversion to poverty, and this is reflected by the dominance properties satisfied by P_α for various ranges of α: monotonicity for $\alpha > 0$, with P_0 just violating it (as increased individual poverty does not affect P_0); the transfer axiom for $\alpha > 1$, with P_1 just violating it; and transfer sensitivity for $\alpha > 2$, with P_2 just violating it. This also signals the versatility of the FGT class in that one can select a specific measure with a desired degree of sensitivity to each of the underlying forms of transformation, namely, increased individual poverty, increased inequality among the poor, or increased inequality among the poor matched by decreased inequality at higher poor incomes.

The final category of axioms draws on the intuition that subgroups of populations can also be viewed as populations having their own levels of poverty—and that these levels should bear some relationship to the poverty level of the overall population. The additive decomposability axiom outlined in the 1984 paper requires overall poverty to be a population share weighted average of subgroup poverty levels. The origins of this axiom are found in the work of Anand [3] in poverty and Theil [127], Bourguignon [16], Shorrocks [120], and Cowell [33] in inequality—although inequality decompositions have an additional between-group term. Measures satisfying additive decomposability have greater applicability—allowing consistent

---

leaves it non-poor; and for continuity, it is invariance to whether a limiting process takes place before or after the measure has been applied.

[12] In its simplest form, monotonicity requires poverty to rise if a poor person's income falls; the transfer axiom says this is also true even if the decrement is matched by a same sized increment to a richer poor person, hence is a "regressive" transfer; transfer sensitivity says this is true even if the regressive transfer among the poor is matched by a "same sized" progressive transfer among the poor that is higher up the distribution, resulting in a "composite" transfer. The three axioms are associated with the three orders of stochastic dominance.

profiles of poverty to be constructed, identifying the characteristics or factors that contribute to poverty, and estimating the contribution of a subgroup to overall poverty.

A second axiom introduced in the paper, called subgroup monotonicity there, but now more commonly known as subgroup consistency, grapples with the fundamental link between subgroup and overall poverty without specifying the structure of that link. It requires overall poverty in a population to rise whenever (i) poverty in a subgroup increases, (ii) poverty in the rest of the population remains the same, and (iii) there is no migration across subgroups. This consistency property gets to the heart of a central policy relevant requirement: that successful regional or targeted poverty policies should, in fact, improve *overall* poverty. The Sen index and other poverty measures that are more sensitive to context and relative position will, in certain cases, violate subgroup consistency. All of the FGT measures satisfy both additive decomposability and subgroup consistency.

2.3 Application

The final section of the paper provided an example using data from the 1970 Nairobi Household Survey to illustrate the use of P_2 in evaluating overall poverty and constructing a profile of poverty. The example was brief with a table listing the poverty levels for subgroups identified by the length of time lived in Nairobi, and showed the subgroup poverty levels and percentage contributions to total poverty. The working paper [53] offered a much more extensive analysis with subgroups defined by several other characteristics and six values of α. Additional figures depicted how rankings among groups are altered as poverty aversion is raised. The longer version also echoed Anand [3] in noting that the values of P_α obtained for $\alpha \geq 2$ "…will not become particularly meaningful until more studies are done which allow comparisons." As noted in Section 3.3.1, examples were soon available, beginning with a series of follow-up studies by Greer [67] and Greer and Thorbecke [68, 69, 70] on food poverty in Kenya. These studies provided key empirical demonstrations of the FGT methodology and, in fact, grew out of the empirical project that helped motivate the original paper.[13]

3 Subsequent Developments

In the 25 years since the publication of the FGT paper, an impressive body of research has arisen making use of its results. Much of this research has applied the measurement techniques to assess trends in poverty or identify characteristics associated with being poor. Several lines of research have evaluated policy options either theoretically or empirically using the measures. A number of studies have added to the pool of measurement techniques by adapting or extending the FGT measures. Others have focused on the axiomatic foundations of poverty measurement or proposed a characterization of the FGT class. This section surveys some of the key results from the subsequent literature. We do not aim to be comprehensive nor particularly unbiased in our presentation; instead, we present a few representative studies from the many that we find interesting.

3.1 Measurement

The FGT paper contributed to poverty measurement by developing a parametric class of measures having desirable characteristics and a simple structure that policymakers could understand. Many authors have continued along this line of research by proposing other evaluation tools including statistical tests for interpreting estimates. A dominance approach to poverty has been developed to understand when poverty comparisons agree over a range of

---

[13] Thorbecke and Greer had been invited by the ILO to analyze poverty in Kenya using the 1975 Integrated Survey. This dataset allows caloric shortfalls and, hence, poverty measures beyond the headcount ratio to be calculated.

poverty lines, or when many different poverty measures agree.[14] With the advent of Sen's capability approach, the literature has constructed measures of multidimensional poverty based on the FGT class; the related measures of chronic poverty and vulnerability are easily obtained by considering time and risk. Finally, there are some interesting applications of the FGT measures to other non-poverty measurement domains. These varied contributions to measurement are now discussed.

3.1.1 Direct extensions

Sen [119] has noted that the choice of a measure depends on the purpose for which it is to be used and, accordingly, several modifications of the basic FGT formula have be proposed to address certain questions. Foster and Shorrocks [65] consider the absolute FGT poverty measure $z^{\alpha}P_{\alpha}$ (used in 4.1.2 below). Rodgers and Rodgers [109] obtain a measure of poverty intensity in a subgroup by dividing a subgroup's P_{α} level by the overall P_{α} level. Bourguignon and Fields [19] propose the discontinuous measure $\delta P_0 + P_{\alpha}$ for $\delta > 0$ and $\alpha > 1$, to obtain more flexible allocative properties than the usual P_{α} class for use in public finance exercises. Ray [107] uses the FGT measures to derive a metric to evaluate the efficiency of transfer systems in reaching their intended targets.

Absolute number versus proportion of poor

The FGT measures, like all poverty measures satisfying replication invariance, are normalized by the total size of the population and therefore do not themselves provide

---

[14] Many of these techniques are described in the book *Poverty Comparisons* by Ravallion [103] whose use of the terms "poverty incidence, depth and severity" for the three main FGT measures P_0, P_1, and P_2, also became common parlance at the World Bank and among policymakers. Other accessible presentations are found in Deaton [40], Foster and Sen [61], Zheng [139], and Duclos and Araar [46].

information on the absolute number of poor.[15] In certain contexts, however, the total number of poor and their aggregate (not average) conditions may be exactly what is important. This observation prompted Kanbur [84, p. 228] to pose the following question: "If the total number of poor goes up but, because of population growth, the percentage of the poor in the total population goes down, has poverty gone up or down?" To take a case in point, in sub-Saharan Africa the headcount ratio corresponding to the $1.25 a day poverty line fell from 53.4% in 1981 to 50.9% in 2005, while the absolute number of poor almost doubled from 212 million to 388 million over the same period. The economist's instinct, influenced by the replication invariance axiom that undergirds the FGT and other standard measures of poverty (and inequality and wellbeing), is to state that poverty has gone down. In contrast, those who work directly with the poor and are burdened by rising numbers may argue that poverty has instead gone up. A simple solution, of course, is to "denormalize" the FGT (or multiply it by the population size) to obtain a measure that is "absolute" in this respect (Foster [56]). Chakravarty, Kanbur, and Mukherjee [25] provide an extension of the FGT index that goes further: it has a parameter that reflects the relative weights to be given to "absolute numbers" versus the "fraction in poverty."

Premature mortality and poverty

The case of sub-Saharan Africa suggests a second, more substantial, reason for extending the FGT class. As emphasized in Section 3.1.4, the FGT measures are static (using data from a single time period) and unidimensional (using data on a single indicator of wellbeing), which does not make them very sensitive to major demographic trends, such as the HIV/AIDS epidemic, or more generally to the positive correlation between premature mortality and poverty that may be influencing the data used by the measures. Kanbur [84, p. 228] raises the premature mortality issue in the following terms: "Suppose the incidence of poverty (and/or the number of poor) goes down because the poor die at a faster rate than the non-poor. Is this a legitimate

[15] That is, unless total population size is known.

‘decrease’ in poverty?” Kanbur and Mukerjee [87] argue that it is not conceptually correct for differentially higher mortality rates among the poor to reduce measured poverty. Indeed, regarding the death of a poor individual due to malnutrition occasioned by low income and lack of assets as an improvement in poverty is at the very least counterintuitive. Similarly, should the increase in the measured level of static, unidimensional poverty that would arise as a result of a successful public health campaign, be viewed as an increase in “actual” poverty? This is an important critique of the entire enterprise of measuring poverty as a function of current income shortfalls—ignoring other dimensions such as health, that are highly complementary to income, and time. However, the problem is even more intractable for two reasons: it involves the transition of death, which is not explicitly addressed by usual assessment methods; it requires an empirical determination of causality and a counterfactual state, which can be both challenging and controversial.

Kanbur and Mukherjee [87] modify the FGT measure to make it sensitive to premature mortality. They specify a normative lifetime close to the range observed in rich countries today. If a poor person dies before this normative age of death, she continues to “live” in the data, with the expected low level of income, until reaching the normative lifetime age. Consequently, the measured level of poverty does not automatically register a decrease in poverty as a result of the premature death. The extension is conceptually attractive and addresses a troubling issue, but has some major practical difficulties if it is to be implemented. Estimating how many years of an individual’s life were “lost” as a consequence of poverty is quite problematic. Estimating the counterfactual income profile of the prematurely dead individual between the actual time of death and the postulated normative lifetime is clearly difficult. And yet both steps are needed for the solution they propose.[16]

---

[16] The discussion is reminiscent of the classic difference between inequality, per se, and inequality due to some other factor (such as a socioeconomic indicator), a distinction that is particularly salient in the context of health. Usually the other factor acts to increase inequality;

Poverty decomposition: Growth versus distribution

The FGT measures have also been used to define and quantify other concepts related to poverty. One methodological issue of great interest to policymakers relates to the quantification of the relative contribution of growth versus redistribution to observed changes in poverty. Datt and Ravallion [38] provide formulae for each of the three main FGT poverty measures that allow an overall change in poverty to be broken down into these two effects (plus a residual term). Their estimates for India and Brazil in the 1980s suggest that the structure of growth was very different in the two countries. A related concept is "pro-poor growth" which gauges the extent to which growth in the mean income is reaching the poor. Son and Kakwani [123], for instance, propose an indicator called the "poverty equivalent growth rate" defined as the rate of growth that would have resulted in the same level of poverty reduction as the present growth rate if the growth process had not been accompanied by a change in inequality. They show that it is a product of the usual growth rate and a pro poor growth index, and compute the rate for the three main FGT measures. Other approaches to pro-poor growth include Foster and Szekely [66] and Duclos [45].

Household size, intrahousehold allocation, and poverty

Poverty levels should reflect real differences in need arising from variations in the size and composition of the household, as well as the inequalities that may be present in the intrahousehold allocation of resources. Empirical applications of the FGT measures typically use surveys providing information on resources at the household level rather than at the individual level. An *equivalence scale* is often used to convert the household resource level to an *equivalent income* for each household member, accounting for the size and makeup of the household. The

---

but in the present context, the disproportionate mortality among the poor serves to lower poverty. Kanbur and Mukerjee [87] would like to evaluate poverty net of the mortality impact.

simplest example assumes that resources are divided evenly among household members and uses the *household per capita income* as an equivalent income. Other equivalence scales account for differences in need across household members (e.g., an infant requires fewer calories than an adult) and the economies of scale as households become larger. The actual allocation within the household may further reflect differences in power between the genders. Specialized datasets are needed to account for this source of variation across households.

The poverty implications of different equivalence scales and of unequal intrahousehold resource allocations have been addressed in the literature using the FGT measures. Coulter et al. [32] explore how estimates of poverty (and inequality) vary with different equivalence scales. Larger households are assumed to have greater needs than smaller ones as represented by a scale parameter (θ) that varies between 0 and 1 and reflects greater needs of larger households compared to a single person household. The authors find that for a given income distribution and most poverty measures (including FGT), poverty first falls and then rises as θ increases from its minimum level. Lanjouw and Ravallion [92] also address this issue, starting with the universal observation that larger households appear to be poorer than smaller households on a per capita basis. The existence of size economies in household consumption cautions against concluding that larger families actually *are* poorer. Using FGT measures and data from Pakistan, the authors find a critical value of the household-size elasticity of the cost of living at which the relationship between poverty and household size switches.

Haddad and Kanbur [72] raise the question of how much difference the existence of intrahousehold inequality could make to conventional measures of inequality and poverty. They investigate this issue using the FGT measures and show that ignoring intrahousehold inequality leads to an understatement of P_1 and an even larger understatement of P_2. However, the situation is reversed for $P_{0,}$ Using the headcount ratio to estimate the magnitude of poverty in the presence of unequal intrahousehold allocation leads to a substantial overstatement of poverty.

3.1.2 Dominance

The main purpose of the FGT paper was to address the aggregation step in poverty measurement—not the identification step. And yet the measures presented there turned out to be a key component in a strategy to avoid the inherent arbitrariness in the identification step by using stochastic dominance. The idea behind the approach of Foster and Shorrocks [62, 63] is simple: in some cases it may be possible to make unambiguous determinations of poverty with respect to the choice of a poverty line; in others, the judgment may depend on the specific line chosen. One can systematically study the former situation to derive the underlying (variable line) poverty ordering for a given poverty measure.

Foster and Shorrocks focus on the three main FGT measures and find that the associated orderings are, in fact, the three orders of stochastic dominance developed in risk analysis, with P_0 yielding first order dominance, P_1 yielding second order dominance (used by Atkinson [7]), and P_2 yielding third order dominance. They note that the three poverty orderings are linked to welfare orderings associated respectively with three classes of additive welfare functions: those with positive marginal welfare of income (or efficiency preferring); those also with diminishing marginal welfare from income (or equity preferring); and those with convex marginal welfare (or transfer sensitive). Consequently, for example, unambiguously lower poverty according to P_1 is equivalent to higher welfare for the Atkinson equity preferring class; and, in the fixed mean case, this implies that the Lorenz curve will be lower (and vice versa).[17] They also observe that the three poverty orderings are nested (in that the poverty ordering for P_0 implies the poverty ordering for P_1, and it in turn implies the poverty ordering for P_2) and obtain analogous results for restricted ranges of poverty lines below some upper threshold.

---

[17] Foster and Shorrocks [64] emphasize the value of the partial ordering approach in unifying the literatures on poverty, inequality and welfare measurement. See also Zheng [139] and Duclos and Makdissi [47].

Atkinson [8] began with these results and made the following interesting observation for continuous additive poverty measures: if the first order dominance holds over a restricted range of lines, then not only can an unambiguous comparison be made for P_0, P_1, and P_2, it must hold for *all* measures satisfying monotonicity. This means that the variable *line* poverty orderings are also variable *measure* poverty orderings. Second order dominance likewise ensures that all the measures satisfying monotonicity and the transfer axiom will agree. This surprising result broadened the reach of the FGT measures, since their variable line poverty orderings were enough to ensure dominance for entire classes of poverty measures. Ravallion [103] saw the potential power of these results for real-world analyses of poverty, and provided a guidebook for practitioners to apply and interpret dominance techniques. He coined the terms "poverty incidence curve," "poverty deficit curve," and "poverty severity curve" for the curves that depict the level of P_α poverty as a function of the poverty line, for $\alpha = 1$, 2, and 3. Jenkins and Lambert [76] recast the exercise in terms of poverty gaps and construct analogous diagrams that neatly illustrate the dimensions of poverty and indicate dominance. For a review of the subsequent literature on poverty orderings, see Zheng [139].

3.1.3 Statistical tools

When measuring poverty in a world of imperfect data, it is useful to formulate statistical tests in order to gauge the confidence of a given comparison. The original FGT paper did not present the associated tools, but since then the literature has provided a steady stream of inference-based research for poverty estimation. It exploits the fact that the FGT measures take the form of a simple mean of a function of incomes given the poverty line. One line of research assumes that the poverty line is fixed and given. Kakwani [78] evaluated the distribution of the estimated poverty values and showed that it asymptotically follows a normal distribution. This allows poverty levels to be compared using a simple difference of means test. Xu [136] suggests an alternative route that shows that the sample counterparts to the FGT and other decomposable measures can be represented as a "U-statistic" (or some function of a U-statistic). Dia [43] first estimates the distribution of income using kernel estimation techniques and then applies a poverty measure to obtain an estimated value.

A second line of research allows the poverty line to be measured with error. Ravallion [103] provides an estimation technique that applies when an absolute poverty line is itself measured with some error, as might occur if the line is itself estimated using, say, a basic needs approach. He provides the distribution of the test statistics appropriate for comparing poverty levels in this situation. Zheng [140] discusses the asymptotic properties of measures of poverty when the poverty line is endogenous. The proposed estimate depends on the underlying population income distribution, and he uses a kernel estimation technique to arrive at the distribution of income based on the incomes of individuals in the sample. Complementing these approaches are tests that explicitly account for the sampling design of survey data. Kakwani [79], Cowell and Victoria-Feser [37], and Zheng [141] among others, have modified the estimation techniques to account for the difficulties that stem from using sample data rather than census data.

One of the most active areas for statistical research associated with the FGT measures concerns tests for poverty orderings or, equivalently, stochastic dominance. Many papers in this literature utilize an FGT measure or some transformation to characterize the dominance criteria. Key papers include Anderson [5], Davidson and Duclos [39], and Barrett and Donald [13], each of whom was clearly motivated by the connection of dominance to the FGT measures. The tests involve the estimation and comparison of population parameters; they essentially use an FGT measure as the sample counterpart for estimation purposes. There are several modules available for standard statistical packages that provide test statistics for the FGT and the associated dominance conditions under various conditions.[18]

---

[18] See, for example, the comprehensive software for distributive analysis (DAD) of Duclos and Araar [46], which highlights the FGT measures and their dominance techniques.

3.1.4 Multidimensional poverty

The FGT measures' axiomatic properties, and additive decomposability in particular, make it a useful instrument for extending poverty measurement beyond its traditional bounds. One key direction is to incorporate other dimensions, such as health, education, and nutrition, into the definition of poverty. A second is to include observations from many time periods to capture dynamic aspects of poverty and to discern between the chronically poor and the transient. A third is to incorporate risk into the discussion, noting that many persons who are not poor are nevertheless vulnerable. Each of these directions takes a broader, multidimensional view of poverty and, as we shall see, the FGT methodology is repeatedly used to construct appropriate measures.

Multiple domains

The original FGT measures were unidimensional, relying on shortfalls in a single variable to reflect poverty. Yet the extent to which a person is well off or poor may well depend on achievements in several distinct and independently important dimensions and, if so, this calls for the development of new multidimensional measures that reflect the complex nature of wellbeing and poverty. Sen's capability framework provides the most comprehensive starting point.[19] In this framework, wellbeing depends on "capabilities" or one's freedom to achieve certain valuable "doings and beings" called functionings, and poverty is viewed as capability (or functioning) deprivation, which for measurement purposes might be indicated by levels of achievement that fall short of minimum levels. Under the impetus of this approach, the analysis and measurement of multidimensional poverty has progressed significantly, and in 1998 two significant papers emerged. The first, a mimeo by Brandolini and D'Alessio [20] outlines a range of possible measurement methodologies that could be drawn upon to construct a multidimensional measure of functioning (and mentions a "work in process" by Bourguignon

[19] See Foster and Sen [61] for a concise presentation of the capability approach.

and Chakravarty). The second, by Chakravarty, Mukherjee, and Ranade [26], proposes a multidimensional extension of the FGT class and specifies a set of appropriate axioms. A number of papers succeeded these including: Tsui [134], which draws quite heavily on the subgroup consistency approach of Foster and Shorrocks [65]; Atkinson [9], which contrasts welfare based and "counting" approaches; Chakravarty and D'Ambrosio [24], which measures the related concept of "social exclusion"; a survey by Chakravarty and Silber [27]; and the twin volumes of Kakwani and Silber [80, 81]. Each contains multidimensional extensions of the FGT methodology. Thorbecke [130] provides an intuitive overview of the related literature; here we discuss three papers in greater depth.

Bourguignon and Chakravarty [17] take as a fundamental starting point that "a multidimensional approach to poverty defines poverty as a shortfall from a threshold on each dimension of an individual's well being" (p. 27). They focus on the two-dimensional case, and include persons who are deprived in at least one dimension in their analysis (later called the "union" approach to identifying the poor). They build a class of multidimensional poverty measures that extends the FGT class to many dimensions. The new measures satisfy a number of desirable axioms and are consistent with the possibility of interacting attributes, a feature they discuss at some length. In the case of substitutes, a person with a larger quantity of attribute k experiences a smaller drop in poverty with an increase in attribute j: for example, the reduction in poverty caused by a unit increase in income would be greater for people with very low education levels than for individuals nearer to the cutoff for education. For complements, the opposite is true: the drop in poverty would be larger for individuals endowed with more education. They argue that, in theory, their poverty indices could be generalized to any number of attributes, but note that this would require assuming the same elasticity of substitution between attributes, which may lessen the appeal of their approach. Their measures are sensitive to the correlation between dimensions, but they require cardinal data, which may restrict their use in practice.

Alkire and Foster [2] present a comprehensive methodology that combines a new method of identifying the poor and a multidimensional extension of the FGT class. The poor are identified using two forms of thresholds: "first a cutoff within each dimension to determine whether a person is deprived in that dimension; second a cutoff across dimensions that identifies the poor using a (weighted) count of the dimensions in which a person is deprived." This identification method extends the "union" approach (where poor persons are deprived in one or more dimensions) and the intersection approach (where the poor are deprived in all dimensions), and it can be used with ordinal data. The aggregation step across individuals employs the FGT measures, appropriately adjusted to account for multidimensionality.[20]

Duclos, Sahn and, Younger [48, 49] develop a dominance approach to multidimensional poverty. They extend the concept of a poverty line in one dimension to a poverty frontier in multiple dimensions, which in turn defines the set of cutoff vectors over which dominance is examined. The measure employed in their dominance exercise is essentially a generalization of the FGT index with separate poverty aversion parameters for each dimension. Their methodology is sensitive to the covariance (or association) across dimensions, but cannot be applied when variables are complements.

[20] We note that the search for a multidimensional poverty measure is not purely an academic pursuit. In 2007, the National Council of Evaluation of the Social Development Policy of Mexico (CONEVAL) was mandated and given the responsibility by the Government of Mexico to develop an operational multidimensional poverty measure that could be used to monitor national poverty and allocate central government funds. The associated law specified the dimensions to be incorporated into the measure. In December 2009, CONEVAL announced that it had selected measures based on the Alkire-Foster extension of the FGT indices.

Time, risk, and vulnerability

The original FGT class is static in that it captures poverty at one point (or period) in time and ignores possible fluctuations in consumption across periods.[21] Yet many of the remaining unresolved issues in poverty analysis relate directly or indirectly to the dynamics of poverty. Policymakers are particularly interested in the conditions under which some households remain chronically poor and others are temporarily in poverty, since the types of interventions needed to alleviate the two types of poverty are very different. Appropriate insurance schemes (such as crop insurance) and other consumption-smoothing measures can be effective in reducing temporary poverty, but are likely to be less effective against chronic poverty. Reducing the latter might require significant investments in human and health capital and some redistribution of assets—particularly land.

McCulloch and Calandrino [96] distinguish three types of chronic poverty: (1) mean consumption across time being below the poverty line; (2) a high frequency of being in poverty over some time (or a high probability of being poor); and (3) a high degree of persistence in poverty. The first two papers considered here use (1) to identify the chronically poor and then apply FGT measures to evaluate its aggregate level. Rodgers and Rodgers [110] regard persistent poverty as a state in which income is less than needs in many consecutive years. Their measurement of chronic poverty is based on a comparison of "permanent income" and "permanent needs." Given a multi-year observation period, Rodgers and Rodgers [110] measure an agent's permanent income as "the maximum sustainable annual consumption level that the agent could achieve with his or her actual income stream over the same period if the agent could save and borrow at prevailing interest rates." In turn, transitory poverty is defined as the difference between annual poverty in a given year and chronic poverty. An agent may be chronically poor but temporarily out of poverty in a given year. Alternatively, an agent who is not chronically poor may experience transitory poverty in some year. The decomposition into the

[21] This subsection draws on Thorbecke [129].

two types of poverty is based on the P_2 (squared poverty gap) measure. In an application to China, Jalan and Ravallion [75] define transient poverty as the contribution of consumption variability over time to expected consumption poverty. The non-transient component is the poverty that remains when inter-temporal variability in consumption has been smoothed out (to the mean consumption level), and this is what they call chronic poverty. Again the measure uses P_2. Foster [58] employs the second definition given above, and regards as chronically poor a person who is in poverty no less than τ share of the time, where τ is a fixed number between 0 and 1. He measures chronic poverty using the FGT measures appropriately adjusted for duration.

The second definition also leads directly to the concept of vulnerability and the extent to which households can protect themselves against a variety of shocks. A number of contributions have attempted to define and operationalize the concept of vulnerability. Christiaensen and Boisvert [29] contrast poverty and vulnerability in the following way. Poverty is concerned with not having enough now, whereas vulnerability is about having a high probability now of suffering a future shortfall. Their notion of vulnerability is the risk of a future shortfall and is expressed as a probability statement regarding the failure to attain a certain threshold of wellbeing in the future. They measure vulnerability as the probability of falling below the poverty line z, multiplied by a conditional probability-weighted function of a shortfall below this poverty line. Consistent with the FGT poverty measure they use a vulnerability-aversion parameter α such that by setting $\alpha > 1$, households with a higher probability of large shortfalls become more vulnerable. A key question at this stage is whether vulnerability and consequent risk-aversion is part and parcel of multidimensional poverty in the sense that certain sets of shortfalls of attributes (deprivations) generate vulnerability or whether vulnerability is a separate dimension of poverty. In a conceptual breakthrough, Ligon and Schechter [94]—also relying on FGT—break down vulnerability into two components reflecting poverty and risk, respectively. The first component is supposed to represent that part of vulnerability due to (chronic) poverty, while the second reflects risk and uncertainty and, presumably, transitory poverty. While this distinction is ingenious and useful in estimating the utility gain that could accrue to the poor, if there were a means to remove their risk-aversion through some social insurance program, it could mask the fact that certain types of current poverty (portfolios of deprivations) render those households more vulnerable. In turn, higher risk by altering the behavior of the poor pushes some of them further into a poverty trap. In this sense, vulnerability (risk) and poverty are inherently inter-related.

Dercon and Krishnan [42] also rely on the FGT measures in exploring seasonal variability in consumption and vulnerability to shocks on a panel data set from different communities in rural Ethiopia. Among their main findings are: 1) a large difference in the levels of consumption within a relatively short period and a large number of poverty transitions (in and out of poverty), apparently linked to seasonal factors; 2) high vulnerability to shocks: i.e., the number of households predicted to fall below the poverty line when serious shocks hit the household and community, based on their seasonal *panel* data set, is about half to three-quarters higher than the poverty estimates obtained using the current *cross-section* estimate in each period. Ignoring seasonal fluctuations tends to lead to a systematic underestimation of poverty.

Stochastic versus structural poverty transitions and an asset-based poverty threshold

Carter and Barrett [21] address a crucial question inherent to the dynamics of poverty, namely, how to distinguish between two very different transitions out of poverty. Individuals may appear to be transitorily poor in a panel study, moving from the poor to the non-poor state over time due to either of two fundamentally different situations. Some may have been initially poor because of bad luck, even though their assets were above the poverty threshold, and their return to the non-poor state simply reflects a return to their normal state (a *stochastic* transition). For others, whose initial assets were below the poverty threshold, the transition may have been *structural*, made possible by the accumulation of new assets or enhanced returns to the assets they already possessed. The authors identify an asset poverty line and derive a family of measures based on the FGT methodology. An important contribution of this asset-based approach to poverty is that it separates those individuals (households) whose assets allow them to move out of poverty from those caught in a low-level equilibrium trap.

3.1.5 Other domains

The FGT measures can be applied to other domains in which there is a well-defined cutoff and shortfalls are considered undesirable. Denny [41], for example, applies the FGT methodology to the domain of education, or more specifically, literacy. He uses test scores from the International Adult Literacy Survey to obtain measures of illiteracy applicable for comparisons across participating countries. The headcount ratio is the usual illiteracy rate; P_1 and P_2 account for the depth and severity of illiteracy. He provides examples where the new measures present a different picture than the literacy rate, and shows how decompositions help clarify the structure of illiteracy within a country. Sahn and Stifel [113] use P_α to measure child malnutrition where the variable is standardized height-for-age z-scores in various countries. Among other results, they find that India has the highest level of malnutrition, whether measured by the traditional incidence, depth, or severity measures. Their analysis includes dominance tests and statistical evaluations. Gundersen [71] applies P_α to the U.S. Department of Agriculture's "food security scale" to obtain the incidence, depth, and severity measures of food insecurity for the U.S.

The FGT measures can also be modified to address the case where observations *above* a cutoff are of concern and the "excess" above the target is used in place of the shortfall below. For example, Basu and Basu [14] apply a modified FGT measure to the distribution of ages in a country to measure the "graying" of the country. Chaplin and Freeman [28] consider measures of the affordability of public housing in the UK based on the distribution of the "rent-to-income" ratios of tenants. The cutoff is set to 25%, and all tenants with rents exceeding 25% of income fail the affordability test. The usual aggregate measure is a headcount ratio of the households failing the test; they select a modified P_3 because of its superior properties and show how the rankings of regions in the UK are altered. Lubrano et al. [95] measure the aggregate productivity of economics departments using a modified P_α measure, where the cutoff is a minimum level of productivity for an individual researcher to be considered active. Jolliffe [77] uses a modified P_α to evaluate the extent to which a population is overweight. The variable here is the distribution of the BMI, and the cutoff is the WHO standard of 25. He argues that the traditional headcount ignores policy-relevant information provided by the "overweight gap" index and the "squared overweight gap" index. Peichl, Schaefer, and Schleicher [102] evaluate the conditions of the rich in Germany using modified FGT measures. Foster, Horowitz and Mendez [60] adapt the FGT measures to assess the aggregate corruption in a region, where a transaction is not considered corrupt until it exceeds a certain cutoff level. These examples illustrate the applicability of the FGT methodology, and show that its impact extends well beyond the domain of poverty measurement.

3.2 Axioms

The axiomatic method played a central role in understanding and communicating the advantages of the FGT class of measures. Many subsequent authors have further investigated the axiomatic structure for poverty measurement, and used the framework for discerning among alternative measures.[22] Additive decomposability and subgroup consistency play central roles in this discussion. Decomposability was motivated primarily by applications and has had a great impact in the empirical and policy uses of poverty measures. The related property of subgroup consistency has revealed itself to be a potent theoretical tool, with links beyond the measurement of poverty to inequality, welfare, and living standards, and even cooperative game theory. In addition, a number of axiomatic characterizations of the FGT class of measures have been presented. In what follows, we discuss a few of these contributions.

---

[22] See, for example, Foster [55], Hagenaars [73], Seidl [114], and Zheng [138].

3.2.1 Decomposability and subgroup consistency

The decomposability property used to motivate the FGT measures is now seen as the standard formula for linking subgroup and overall poverty levels. A lengthy list of empirical applications of decomposable measures now exists—some examples are given below in Section 3.3.1—and a parallel body of work has explored the conceptual implications of the property. In an early discussion of decomposability, Foster [55] stressed that while the property is not needed for every case where subgroup poverty levels are evaluated, it becomes essential when the goal is an assessment of the contribution of subgroup poverty to overall poverty. Several papers restrict attention to general classes of decomposable measures. Atkinson [8], for example, defines an "additively separable" class of poverty measures over continuous income distributions that are also decomposable.[23] Foster and Shorrocks [65] axiomatically derive several general classes of decomposable poverty measures using the allied property of subgroup consistency.

The subgroup consistency axiom, which appeared in the FGT paper in a slightly less demanding form, was originally motivated by examples from inequality analysis found in earlier versions of Cowell [34, 35] and in Mookherjee and Shorrocks [97]. A similar property, with restrictions on subgroup means, leads to the most natural characterization of the generalized entropy measures of inequality by Shorrocks [121]. Foster and Shorrocks [65] explored the implications of subgroup consistency for poverty measurement and showed that subgroup consistent indices are in essence monotonic transformations of additively decomposable measures. These results justify decomposable indices, not just for their empirical usefulness, but for their unique role in ensuring consistency between subgroup and overall levels. More recently, and in a different context, the property has been shown to select Atkinson's [7] parametric class

[23] Atkinson [8] uses an additive form to mirror the structure of the welfare functions in Atkinson [7] and does not mention decomposability.

of equally distributed equivalent income (ede) functions and the other "general means" from among all possible income standards.[24]

A number of arguments have been advanced against decomposability and subgroup consistency as universal requirements for poverty measures. Sen [116], for example, argues against decomposability in part because of its unlimited applicability to all partitions of the population, whether salient or not.[25] Instead, he proposes a more limited application of decomposability to some partitions (where linkages are within subgroups) and not others (where linkages are across subgroups). Sen's argument is based on a rejection of symmetric unidimensional poverty measurement, and hence is perhaps better seen as a call for multidimensional measurement than as a rejection of decomposability, per se. Foster and Sen [61] provide a balanced view of the arguments for and against subgroup consistency and decomposability, and note that the presence of a poverty standard makes the axioms more plausible for poverty measures than for inequality measures.

3.2.2 Characterizations

A first approach to characterizing the FGT can be found in Basu and Basu [14], who modify Sen's three-step procedure to derive their variation of the FGT indices. Foster and Shorrocks [65] assume subgroup consistency and derive the class of relative measures (satisfying scale invariance) and the class of absolute measures (satisfying translation invariance).[26] They

---

[24] An income standard summarizes an entire distribution as a single representative income level and satisfies various invariance axioms, a normalization axiom, and linear homogeneity. Functions of this sort underlie nearly all inequality and poverty measures; see Foster [57], [58].

[25] Despite Sen's disagreement with decomposability, one of the first papers to cite FGT was Sen [115].

[26] See Donaldson and Weymark [44] for a discussion of absolute measures. The absolute FGT measures, defined by $z^\alpha P_\alpha$, are routinely employed in tests of stochastic dominance.

show that the intersection only includes the headcount ratio or certain transformations (see, also, Zheng [137], who drops subgroup consistency). They then ask if there might be pairs of measures, one relative and one absolute, that are compatible in that they render the same judgments when the poverty line is fixed. They prove that (P_α, $z^\alpha P_\alpha$) is essentially the only such pair, where the latter is the absolute FGT measure.

Ebert and Moyes [50] extend these results to the orderings represented by poverty measures at a given poverty line (each ordering does not make comparisons across poverty lines). They convert the two invariance axioms into ordinal forms, and note that the FGT orderings generated by P_α (and $z^\alpha P_\alpha$) are consistent with both axioms (indeed they are the same ordering). Their main theorem shows that the FGT orderings are essentially the only ones consistent with the two axioms. Ebert and Moyes also introduce a new concept analogous to an income standard or ede in inequality analysis. The equivalent societal income (ESI), or the amount of income which, if received by all individuals in the population, will yield the same poverty level as the actual income distribution, is a cardinal representation of the underlying "lower poverty" ordering that could be useful for policy purposes.

Chakraborty, Pattanaik, and Xu [22] characterize measure P_2 with the help of a property they call "equivalent transfer". As noted in Foster, Greer, and Thorbecke [54], P_2 just violates transfer sensitivity, in that a transfer of the same size between two poor persons the same distance apart has the same effect on poverty, no matter the location of the pair. This equivalent transfer axiom states this in the form of a requirement, and their main result shows that given a range of standard properties, this property uniquely identifies P_2.

3.3 Applications

The FGT class has proven to be very useful for evaluating the extent of poverty across space and time. Empirical applications abound, and virtually every country has been analyzed using P_2, P_1, and P_0, at one time or the other. Most of the applications make use of the decomposability property to analyze the important correlates of the incidence, depth, and severity of poverty and set the stage for informed discussion of policies to confront poverty. Accordingly, the FGT class has been instrumental in the design of these policies, in part due to its decomposable structure, but also because of its specific functional form that generates straightforward formula for policy design. The following provides a sample of the broad and varied literature on the applications of the FGT class of measures.

3.3.1 Empirical applications

The first empirical applications of the FGT methodology were undertaken by Greer [67] and Greer and Thorbecke [68, 69, 70] using the Integrated Rural Household Survey data for Kenya. A detailed food poverty profile of Kenyan smallholders (constituting over 70 per cent of the total population) was derived. National food poverty, as measured by P_2, was decomposed according to six sets of characteristics: region of residence, household size and composition, household landholding size, cropping pattern and degree of market involvement, type of employment, and by gender and marital status of the head of the household. Regional food poverty lines were established based on the diet composition actually consumed by the poor reflecting the different regional tastes and preferences and local food prices. These studies illustrated and demonstrated the decomposable property of the P_2 measure based on population-share weights and set the stage for a myriad of subsequent poverty profiles based on the FGT methodology worldwide.

The literature contains hundreds of empirical applications using the FGT methodology, and we will not attempt a survey here. Instead, we highlight briefly some selective examples of sectoral and regional poverty decompositions.

Sectoral decompositions

The sectoral pattern of growth is a key determinant of the impact of growth on poverty. Thus, in an early application of all three FGT measures, Huppi and Ravallion [74] examine the structure of poverty in Indonesia by sector of employment, and how it changed during the adjustment period 1984 to 1987. They found that gains to the rural sector in key regions were quantitatively important to Indonesia's success in alleviating poverty. Most poverty exists—and most gains in alleviating poverty were made—in the rural farming sector. These gains were associated with crop diversification and continued growth in off-farm employment. The government's adjustment program favored rural areas and were crucial to Indonesia's evident success at maintaining momentum in alleviating poverty.

The impact of a sector's output on poverty alleviation can be direct through the increase in incomes accruing to the poor households who contributed through their labor or land or other resources to the sector's growth of output. But another part of poverty reduction results from the interdependence of economic activities (the socioeconomic groups' spending and re-spending effects), which can be estimated within a general equilibrium framework such as a Social Accounting Matrix (or SAM). Thorbecke and Jung [132] applied the FGT measure to a relatively highly disaggregated Indonesian SAM (including 75 sectors, 8 household groups and 24 different production activities). They found that total poverty reduction effects originating from agricultural production activities are highest, followed by services and informal activities—a finding that can be generalized to most settings. Among manufactures, food processing and textiles which have closer inter-production linkages with agriculture or are more labor intensive (especially of unskilled labor), made relatively large contributions to poverty alleviation.

Spatial and regional extensions

A critical issue in regional poverty decompositions relates to the choice of poverty lines. For instance, in the light of distinct dietary regimes and differences in relative prices, would it be more appropriate to use regional poverty lines rather than the same national poverty line?

There are currently two main monetary methods of setting the poverty line, i.e., the Cost of Basic Needs (CBN) and the Food-Energy-Intake (FEI) methods. The CBN approach has the advantage of ensuring consistency (treating individuals with the same living standards equally), while the FEI approach has the advantage of specificity reflecting better the actual food consumption behavior of individuals around the caloric threshold given their tastes, preferences, and relative prices.

It has been argued by Ravallion and Bidani [104] that in order to make valid welfare comparisons, the reference basket (bundle) yielding the caloric threshold should remain constant. The monetary poverty line at any point in time is then obtained by multiplying the constant quantitative reference basket by the variable price vector to obtain z at current (nominal) prices and then deflating it by an appropriate price index (often the consumer price index, CPI) to express z in real terms. Given the crucial importance of context-specific conditions in shaping the perception of poverty, it can be argued that the setting of z at a more location-specific level would lead to a more accurate appraisal of poverty. The use of a (normative) national or even provincial poverty line in the light of major intra-regional and inter-village differences in socioeconomic conditions can distort the poverty diagnosis at the local level. Again, this illustrates the inherent conflict between the specificity and consistency criteria. It is not possible to satisfy both simultaneously. The trade-off between welfare consistency (using the same national food basket in all regions) and being realistic and faithful to different regional preferences and diets (which calls for specific regional poverty lines) is unambiguous [129]. The earliest applications of FGT (in particular P_2) to measuring food poverty in Kenya by Greer and Thorbecke adopted the FEI methodology to derive the different food baskets actually consumed by the poor around the caloric RDA and corresponding regional food poverty lines. The welfare consistency criterion was violated in favor of capturing the actual consumption behavior of households in different settings. Subsequently, most attempts at deriving regional poverty estimates relied on the CBN approach on the ground that this method besides being welfare consistent was also more robust. One of the most elaborate critical comparisons of these two methods in a regional application to Mozambique was undertaken by Tarp et al. [126], and it is probably fair to state that the choice of approach is still being debated.

3.3.2 Policy applications

In contrast to the empirical applications described above, which seek to understand the patterns of poverty observed in the real world, the policy applications considered in this section are explicitly normative, with a poverty measure replacing a welfare function as the objective function. The framework has led to a rich literature on the optimal design of policies to minimize P_α under varying assumptions about feasible policy instruments. This, in turn, has helped guide the design and evaluation of a prominent contingent cash transfer program and other development programs.

The idea that P_α could be used in policy design was not explicitly discussed in the FGT paper, although it was apparent that decomposability and subgroup consistency were partially motivated by policy considerations. Much of the literature on optimal design of poverty programs makes use of the specialized form of the P_α class of measures and, in particular, the individual poverty function g_i^α underlying the aggregate measure. Differentiation shows that for $\alpha > 0$, the marginal impact of a transfer of income on i's poverty is proportional to $-g_i^{\alpha-1}$; hence, the marginal impact on P_α of equal sized transfers to all members of a population subgroup is proportional to $-P_{\alpha-1}$. This special property of the FGT class—that the impact on P_α (or g_i^α) depends on $P_{\alpha-1}$ (or $g_i^{\alpha-1}$)—has been exploited repeatedly in the literature that followed the important work of Kanbur [82, 83]. We now briefly discuss this literature.

Consider a world in which a government (or some other entity) is to allocate funds across a population in such a way as to maximally reduce poverty, but is limited by information and policy tools. This is similar to a traditional optimal taxation exercise [85], but with a more focused poverty objective instead of a social welfare function. The simplest case of "perfect targeting" is where the government has full information about consumption (or income) levels and can transfer funds freely to individuals. How would it allocate a limited budget of welfare benefits? Bourguignon and Fields [18] take up this "perfect targeting" scenario and derive the optimal (poverty minimizing) form of direct income transfers from a fixed budget when poverty is measured by the FGT class. They find that for P_α having $0 \leq \alpha < 1$, the optimal policy raises the richest poor out of poverty; for P_1 a best policy transfers funds indiscriminately among the poor, so long as none is lifted above the poverty line; and for P_α having $\alpha > 1$, the optimal policy begins with the poorest person and raises this income level until the next poorest income level is reached, at which point both are raised up, and so forth. This exercise provides an additional reason why the headcount ratio is not a very good objective: it yields a remarkably regressive policy. In contrast, the distribution-sensitive measure P_2 always targets the poorest of the poor.

The above scenario assumes perfect information—and yet information is rarely, if ever, perfect, and is always costly. One can avoid information costs by providing equal benefits to all households regardless of their incomes and characteristics (the no targeting policy) but only at the expense of large leakages that would allow a significant part of the anti-poverty budget to be used in a way that does not directly reduce poverty. An alternative between the extremes of perfect information and none at all is the case where the population can be partitioned into identifiable subgroups. It is assumed that the policymaker can adjust welfare benefit levels across subgroups in response to group attributes but cannot vary benefits within a subgroup. Assuming that there is a fixed budget devoted to the reduction of poverty as measured by P_α for $\alpha > 1$, how should the budget be allocated among the mutually exclusive groups?

One might naturally expect the policymaker to target the subgroup with the highest level of poverty according to P_α in order to minimize P_α. However, Kanbur [82, 83] showed that the optimal budgetary rule would allocate funds first to the subgroup whose $P_{\alpha-1}$ level is highest. A similar optimal policy is obtained in the model of poverty alleviation and region-specific food subsidies of Besley and Kanbur [15], a model that would apply equally well to any commodity consumed by the poor. Ravallion and Chao [105] provided a useful numerical algorithm for this framework, and also discuss the optimal targeting mix across urban and rural areas to minimize P_2. They note that the headcount ratio P_0 in many countries is many times higher in rural than urban areas, but that this does not imply that urban poverty should *not* be targeted. This decision should depend on P_1 levels, which in principle could be higher in urban areas (but as noted by Levy [93], typically are not).

Thorbecke and Berrian [131] considered the optimal budgetary rule in a SAM based computable general equilibrium model that allows for interactions between subgroups, and shows how the rule becomes more complex. In particular, if P_2 is the objective function to be minimized, it may not be optimal to target the group with the highest P_1, due to the differential patterns of spending across subgroups and the subsequent impact of the interactions.

In an influential paper, Levy [93] considered the design of a poverty program for the extremely poor in Mexico. A key part of his discussion concerned the choice of poverty measure:

> Indicators of poverty should incorporate concerns about its severity and distribution; the head-count ratio fails to do this, as do other 'indices of marginalization' computed by government agencies and currently used to identify the poor. The Foster-Greer-Thorbecke poverty index [P_α] satisfies axioms with respect to severity and distribution of poverty, can be separably decomposed, allows measurement of the contribution of each region to total poverty, can serve to rank regions for delivery of benefits, and can be used to monitor progress in poverty alleviation [93, p. 83].

He argued for poverty information to be gathered at the state and county level; the resulting geographical poverty profile should then used to target poverty programs that “exploit the complementarities among nutrition, health and education,” where targeting is done in accordance with Besley and Kanbur [15], and also “could be made contingent on (parents) bringing their children periodically for inoculations and other … medical attention.” [93, p. 63]. This was the template for the Progresa program, which was implemented by Levy under President Zedillo in 1997. This pioneering conditional cash transfer (CCT) program was designed with the goal of minimizing P_2.[27]

More recently, FGT measures have been used by Morduch [98, 99] to understand the poverty impact of different microfinance strategies. Suppose that a self-sustaining institution is

---

[27] The Mexican law implementing *Progresa* required the program to be targeted to P_2 applied to an aggregated variable and, indeed, the formula of P_2 appears in the statement of the law.

only able to target the “richer” poor (say, with incomes at 90% of the poverty line), while a program receiving an external subsidy can target poorer households (say, with 50% of the poverty line). Assume further that both are able to achieve the same (marginal) increase in household income. Then poverty measured by P_0 would be unchanged for both; P_1 poverty would fall equally for both; but P_2 would register a five-fold decrease for the subsidized program as compared to the self-sustaining one (with an even greater differentiation when P_3 is used).[28] Viewing the subsidies as the cost of better targeting, and the associated larger decrease in poverty as the benefit, the net benefit may well be larger for a subsidized microfinance institution than one that is fully sustainable.

Two final examples link the FGT indices and two policy instruments for poverty alleviation. One concerns the Millennium Development Goals (MDGs) of the United Nations, adopted in 2001 by the UN member nations. The first MDG is “to eradicate extreme poverty and hunger”; its first target is “to halve…the proportion of people whose income is less than one dollar a day”; and measures P_0 and P_1 are the first two indicators for monitoring progress of the target. One question concerns the cost of achieving a country-by-country percentage reduction as compared to achieving the same overall reduction in a “poverty efficient” way. Anderson and Waddington [6] consider this question for P_0, P_1, and P_2 under certain simplifying assumptions, and show that for the headcount the country-by-country targets entail much greater poverty reductions for the poorest countries than in the poverty efficient allocation; whereas, for the poverty gap P_1 and squared gap P_2 this difference appears to be much smaller. In contrast, Collier and Dollar [31] show that the existing allocation of aid across countries is quite far from an optimal allocation, and that the optimal allocation appears to be robust to the choice between the three FGT measures P_0, P_1, and P_2.

---

[28] The derivative of the individual poverty function for P_0 is 0; for P_1 it is proportional to -1; and for P_2 it is proportional to $-g_i$.

A second example concerns the Poverty Reduction Strategy Papers (PRSPs), introduced in 1999 by the World Bank and International Monetary Fund, to help low-income countries assess and combat poverty. Constructing a PRSP is particularly challenging for the poorest countries due to the specialized research capabilities that are needed.[29] It therefore was fortunate that over the period 1995–2005 approximately 150 African economists were trained in poverty analysis under the auspices of a large-scale collaborative project initiated by the African Economic Research Consortium.[30] National research teams from 15 countries in sub-Saharan Africa completed poverty profiles and assessments that relied extensively on the FGT methodology. These country case studies typically became the basis for the country PRSPs.

4 Prospective Areas

The technology of poverty measurement continues to improve as it is adapted to address new conceptual, empirical, and policy challenges. Many of the topics discussed above are quite active areas for research. In particular, multidimensional poverty measurement has risen in importance due in part to Sen's capability approach and increased demand by individual countries and international organizations. The statistical literature on dominance is also expanding as are applications to poverty evaluation and other fields relying on this technology. Policy and empirical applications are ongoing in many countries and sectors. What will be the next direction for expanding this useful technology? Where will the literature take us over the next 25 years? We conclude our paper with a brief outline of some possible directions for future work. We do not attempt to provide a definitive response, but rather suggest a number of topics that are especially of interest to us.

---

[29] For example, the guide to preparing a PRSP includes a heavy dose of FGT measurement technology, and even summarizes the Bourguignon and Fields [18] results.

[30] This project was co-directed by Erik Thorbecke.

4.1 Qualitative data

As poverty measurement is taken to other domains or extended into a multidimensional environment, a crucial emerging issue is how to measure poverty when data do not have the characteristics of income, which is typically taken to be cardinal and comparable across persons. How is poverty or deprivation to be measured in the presence of ordinal, categorical, or qualitative data?[31] Must we retreat to a headcount ratio, or can we continue to evaluate the depth or distribution of deprivations—key benefits provided by the higher order FGT measures when the variable is cardinal? This issue can also arise in unidimensional studies but is almost inevitable in discussions of multidimensional poverty where data on capabilities and functionings can have the most rudimentary of measurement characteristics.

4.2 Axioms versus intuition

The original Sen critique of the headcount ratio and the income gap ratio motivated the development of an array of distribution sensitive poverty measures, including the squared poverty gap measure P_2. These measures satisfy additional axioms and thus are deemed to be superior; but in moving from "partial" indices (cf, Foster and Sen [61], p. 168) to aggregate measures, has something been lost? Partial index values are often inherently meaningful and convey information on an important aspect of poverty (e.g., $H = 0.40$ indicates that 40% of the population is poor). In contrast, the numerical values of distribution sensitive measures typically become meaningful when compared to other values (as emphasized by Anand [3] and Foster, Greer, and Thorbecke [53]). The numerical values of P_2, which are often close to zero due to the squaring of the normalized gaps, can be difficult to interpret even though the formula is clear and even transparent. And this problem is even greater for P_α with $\alpha > 2$, which satisfy an additional

[31] See for example the discussion in Alkire and Foster [2] or Kanbur [86].

axiom—transfer sensitivity—and yet are rarely if ever used.[32] Is there a trade-off between axiomatic desirability and the information conveyed by a poverty measure? How can this tradeoff be evaluated and when should the balance tilt in favor of axioms or, alternatively, in favor of simplicity?

It should be remembered that the headcount ratio is *still* the most commonly used measure of poverty; it would be interesting to understand the factors behind its prevalence despite its failure of fundamental criteria for measuring poverty. Alternatively, might there be ways of making the numerical values of P_α more transparent—especially for larger α? One approach that is inherent in Ravallion [103] is to use the three measures P_0, P_1, and P_2 in tandem to help clarify the additional information each brings to the table. The poverty incidence measure gives information on frequency, the poverty depth measure adds information on average gaps, while the poverty severity index includes information on the distribution by concentrating on the poorest poor. Another possibility suggested by Subramanian [125] and Foster [57] is to consider the subgroup consistent transformation $(P_\alpha)^{1/\alpha}$ for $\alpha > 0$ that is not additively decomposable, but has values that are easier to interpret.[33] This is analogous to the tradeoffs between the decomposable generalized entropy inequality measure (whose values are difficult to interpret) and the Atkinson class of inequality measures (whose values have an intuitive interpretation via

---

[32] For example, consider the following quote from the IDRC website on poverty measurement: "Notwithstanding the above, interpreting the numerical value of FGT indices for α different from 0 and 1 can be problematic. We can easily understand what is meant by a proportion of the population in poverty or by an average poverty gap, but what, for instance, can a squared-poverty-gap index actually signify? And how to explain it to a government Minister?..." See http://www.idrc.ca/en/ev-103707-201-1-DO_TOPIC.html.

[33] As a general mean of normalized gaps it measures poverty in the space of normalized gaps, and has higher values at higher α. It is related to the second measure in Clark, Hemming and Ulph [30] and to the ESI of Ebert and Moyes [50].

welfare), which is not decomposable. It would be very useful to explore the tension between axiomatic acceptability and meaningfulness in the context of poverty measurement.

4.3 Relativities and absolutes

As observed by Foster [56], there are many competing notions of “absolute” and “relative” in the measurement of poverty. For example, P_2 is relative in terms of its invariance property and its treatment of population sizes; it typically uses an absolute poverty line that does not adjust when incomes rise. It would be useful to subject the now traditional assumptions to a more systematic analysis, and to explore alternatives that may be more appropriate in certain circumstances. The axiom of scale invariance, for example, might be the subject of additional scrutiny. For example, this axiom is only applicable when analyzing distributions at different poverty lines, and yet one might argue that only common line comparisons are meaningful. At the same time, the sharp divide between relative poverty lines, which are typically used in the more developed countries, and absolute poverty lines, which are often used in developing countries, needs to be bridged. How this is to be done, and how the resulting cutoff will interact with the poverty measure, are important issues for investigation. Some progress in this direction includes work by Ali and Thorbecke [1], Atkinson and Bourguignon [10], Foster and Szekely [66], and Ravallion and Chen [106]; for a treatment of a related line of research in the measurement of welfare, see Atkinson and Brandolini [11].

As noted above, the AIDS epidemic has brought to the fore a related difficulty: if a disease disproportionally affects poorer persons, and through death removes them from consideration by the measure, then, indeed, average poverty may well be seen to diminish. Who is to be taken into account when measuring poverty? Over what time frame? This problem is certainly related to the multidimensional aspect of poverty (including the problem of time), but it can also be viewed as an example of a broader category of fundamental issues concerning the “denominator” of poverty measurement: who is to be included in the calculation of the poverty value? The complexity of the relative/absolute divide expands in the face of a multidimensional approach to poverty and is related to the issue of how to weigh deprivations in one dimension against the deprivations in another. Could one dimension be measured against a relative cutoff while a second used an absolute one? What might this imply for the trade off between dimensions as the relative cutoff changes? There are many questions that need to be addressed.

4.4 Identification versus aggregation

Sen's [118] observation that the aggregation step is an important component of poverty analysis has led to the development of many alternative aggregation methods. The identification step, by contrast, has changed very little in structure from the time that Rowntree [111] constructed a poverty line of just over a pound a week to identify poor families in York. Given the achievements in aggregation, it may be an appropriate time to re-evaluate the notion of a poverty line and its role in identifying and targeting the poor. Some have suggested that the entire construct of a poverty line should be thrown out.[34] This is an intriguing possibility, which if accepted would lead to a radical overhaul of the way poverty is envisioned and would, in turn, provoke many questions. If the identification step is to be dropped, what should replace it? If it is to be retained, how can it be accomplished without a poverty line? How is targeting to be defined and implemented? Some authors criticize the abrupt 0-1 nature of the poverty line and replace it with a fuzzy approach. Is this a good solution or does it simply multiply the arbitrariness? If the identification step is altered, what then is to become of the focus axiom as a property of aggregate poverty measures? Foster and Szekely [66] use an inequality averse general mean or ede as a "poor income standard" having no poverty line but undeniably emphasizing the poorest incomes. Is this an effective way of combining concerns for poverty and inequality in one

[34] Deaton [40, p. 144], for example, expresses deep skepticism about poverty lines: "…I see few advantages in trying to set a sharp line, below which people count and above which they do not." However, he also notes the practical advantages of poverty lines and continues to use them and the FGT measures in his empirical work.

measure, or is there a superior approach?[35] Further work is needed to evaluate practical alternatives to poverty line identification.

Second, although there is a perception that the identification step is straightforward in the unidimensional case, some interesting and subtle points remain concerning the link between the identification step, the poverty line, and the aggregation step. Imagine a methodology that used one poverty line to identify the poor and a second, higher poverty line as the standard against which to aggregate. This might be quite natural if the goal were to determine the conditions of the poorest poor as a separate group within the larger set of the poor.[36] Of course, this would not matter if the measure were the headcount ratio, since it relies only on the frequency of the poor, not how far they fall below the poverty line. However, the values given by P_1 and P_2 can be greatly altered by having a different line for identification and aggregation. What rules out such a disconnect between the two steps of poverty measurement? Can axioms be reformulated to apply to the overall methodology of poverty measurement, including the identification and aggregation step? This is a subtle problem, but one that deserves additional thought.

Third, the importance of the identification grows much larger when poverty is taken to be a function of many multiple dimensions. Now the set of the poor can assume a variety of shapes including the two extremes of the union and the intersection and all possibilities in between. Which form of identification is the most natural? Can axioms be devised for identification methods alone? How should axioms for measures be reformulated to apply to the entire methodology, including the identification and aggregation steps? Some initial steps have been taken by Alkire and Foster [2], but much remains to be done.

---

[35] Fields [52] considers other ways of combining inequality and poverty measures.

[36] See Levy [93]. The integrated measurement of the "ultra-poor" is an important area for future research.

5 Conclusions

This paper has revisited the literature associated with the FGT class of decomposable poverty measures. The original paper [54] introduced the new class of poverty measures, helped to interpret and justify them using axioms, and provided a tangible illustration of the methods based on Kenyan data. Since then, in addition to a large number of empirical applications worldwide, the FGT class has been extended conceptually and enriched in many domains such as poverty dominance, the interrelationship between poverty and inequality, the dynamics of poverty, vulnerability, the measurement of multidimensional poverty, and as a social objective function, to name only a few. Opinions may vary as to whether the FGT class is uniquely well suited to measure poverty, or is one of a range of acceptable forms.[37] However, we hope that we have convinced the reader that the class has played a key role in this literature since it was introduced some twenty-five years ago.

References

1. Ali, A., Thorbecke, E.: The state and path of poverty in Sub-Saharan Africa: some preliminary results. J. African Economies **9(1),** 9–40 (2000)
2. Alkire, S., Foster, J.: Counting and multidimensional poverty measurement. OPHI Working Paper 7, University of Oxford (2007)
3. Anand, S.: Aspects of poverty in Malaysia. Rev. Income Wealth **23**, 1–16 (1977)
4. Anand, S.: Inequality and poverty in Malaysia: measurement and decomposition. Oxford University Press, New York (1983)
5. Anderson, G.: Nonparametric tests of stochastic dominance in income distributions. Econometrica **64(5),** 1183–1193 (1996)

---

[37] See Hagenaars [73] and Chakraborty, Pattanaik, and Xu [22] for particularly favorable assessments, and Foster and Shorrocks [65] and Ebert and Moyes [50] for axiomatic characterizations; alternative views are suggested by Foster [55] and Foster and Sen [61].

6. Anderson, E., Waddington, H.: Aid and the MDG poverty target: How much is required and how should it be allocated? Oxford Devel. Stud. **35(1)**, 1–35 (2007)
7. Atkinson, A.B.: On the measurement of inequality. J. Econ. Theory **2**, 244–263 (1970)
8. Atkinson, A.B.: On the measurement of poverty. Econometrica **55(4),** 749–764 (1987)
9. Atkinson, A.B.: Multidimensional deprivation: contrasting social welfare and counting approaches. J. Econ. Inequality **1**, 51–65 (2003)
10. Atkinson, A.B., Bourguignon, F.: Poverty and inclusion from a world perspective. In Stiglitz, J.E., Muet, P.-A. (eds.) Governance, Equity, and Global Markets: the Annual Bank Conference on Development Economics — Europe, pp. 151–166. Oxford University Press, Oxford (2001)
11. Atkinson, A.B., Brandolini, A.: On analysing the world distribution of income. Working Papers 97, ECINEQ, Society for the Study of Economic Inequality (2008)
12. Barrett, C.B. (ed.): Development Economics: Critical Concepts in Development Studies. Routledge, New York (2007)
13. Barrett, G.F., Donald, S.G.: Consistent tests for stochastic dominance. Econometrica **71**, 71–104 (2003)
14. Basu, K., Basu, A.: The greying of populations: concepts and measurement. Demography India **16**, 79–89 (1987)
15. Besley, T., Kanbur, R.: Food subsidies and poverty alleviation. Econ. J. **98**, 701–771 (1988)
16. Bourguignon, F.: Decomposable income inequality measures. Econometrica **47**, 901–920 (1979)
17. Bourguignon, F., Chakravarty, S.: The measurement of multidimensional poverty. J. Econ. Inequality **1(1),** 25–49 (2003)
18. Bourguignon, F., Fields, G.: Poverty measures and anti-poverty policy. Rech. Econ. Louvain **56(3-4),** 409–427 (1990)

19. Bourguignon, F., Fields, G.: Discontinuous losses from poverty, generalized P_α measures, and optimal transfers to the poor. J. Public Econ. **62(2),** 155–175 (1997)

20. Brandolini, A., D'Alessio, G.: Measuring well-being in the functioning space. Mimeo, Banca d'Italia, Rome, Italy (1998). In: Chiappero-Martinetti, E. (ed.) Debating Global Society: Reach and Limits of the Capability Approach, pp. 91–156. Feltrinelli Foundation, Milan (2009)

21. Carter, M.R., Barrett, C.B.: The economics of poverty traps and persistent poverty: an asset-based approach. J. Devel. Stud. **42(2)**, 178–199 (2006)

22. Chakraborty, A., Pattanaik, P.K., Xu, Y.: On the mean of squared deprivation gaps. Econ. Theory **34**, 181–187 (2008)

23. Chakravarty, S.R.: A new index of poverty. Math. Soc. Sci. **6**, 307–313 (1983)

24. Chakravarty, S.R., D'Ambrosio, C: The measurement of social exclusion. Rev. Income Wealth **52(3),** 377–398 (2006)

25. Chakravarty, S.R., Kanbur, R., Mukherjee, D.: Population growth and poverty measurement. (2002). Revised version published in Soc. Choice Welfare **26**, 471–483 (2006)

26. Chakravarty, S.R., Mukherjee, D., Ranade, R.R.: On the family of subgroup and factor decomposable measures of multidimensional poverty. Res. Econ. Inequality **8**, 175–194 (1998)

27. Chakravarty, S.R., Silber, J.: Measuring multidimensional poverty: the axiomatic approach. In N. Kakwani, J. Silber, (eds.) Quantitative Approaches to Multidimensional Poverty Measurement, pp. 192–209. Palgrave Macmillan, New York (2008)

28. Chaplin, R., Freeman, A.: Towards an accurate description of affordability. Urban Stud. **36(11),** 1949–1957 (1999)

29. Christiaensen, L.J., Boisvert, R.N.: On measuring household food vulnerability: case evidence from Northern Mali. WP 2000-05, Dept. of Agricultural, Resource, and Managerial Economics, Cornell University (2000)

30. Clark, S., Hemming, R., Ulph, D.: On indices for the measurement of poverty. Econ, J. **91**, 515–526 (1981)
31. Collier, P., Dollar, D.: Aid allocation and poverty reduction. Europ. Econ. Rev. **46**, 1475–1500 (2002)
32. Coulter, F.A.E., Cowell, F.A., Jenkins, S.P.: Equivalence scale relativities and the extent of inequality and poverty. Econ. J. **102**, 1067–1082 (1992)
33. Cowell, F.A.: On the structure of additive inequality measures. Rev. Econ. Stud. **47**, 521–531 (1980)
34. Cowell, F.A.: The structure of American income inequality. Rev. Income Wealth, **30(3),** 351–375 (1984)
35. Cowell, F.A.: Inequality decomposition: three bad measures. Bull. Econ. Res. **40**, 309–312 (1988)
36. Cowell, F.A.: Measurement of inequality. In Atkinson, A.B., Bourguignon, F.: Handbook of Income Distribution, pp. 87–166. North-Holland, Amsterdam (2000)
37. Cowell, F.A., Victoria-Feser, M.P.: Poverty measurement with contaminated data: a robust approach. Europ. Econ. Rev. **40(9),** 1761–1771 (1996)
38. Datt, G., Ravallion, M.: Growth and redistribution components of changes in poverty measures: a decomposition with applications to Brazil and India in the 1980s. J. Devel. Econ. **38(2),** 275–295 (1992)
39. Davidson, R., Duclos, J.Y.: Statistical inference for stochastic dominance and for the measurement of poverty and inequality. Econometrica **68(6),** 1435–1464 (2000)
40. Deaton, A.: The Analysis of Household Surveys: A Microeconometric Approach to Development Policy. The Johns Hopkins University Press, Baltimore, MD (1997)
41. Denny, K.: New methods for comparing literacy across populations: insights from the measurement of poverty. J. Roy. Statistical Society, Series A–Statistics in Society **165**, 481–493 (2002)

42. Dercon, S., Krishnan, P.: Vulnerability, seasonality and poverty in Ethiopia. In Baulch, B., Hoddinott, J., (eds.) Economic Mobility and Poverty in Developing Countries, pp. 25–53, Frank Cass Publishers, London (2000)
43. Dia, G.: Nonparametric estimation of income distribution and poverty index. Comptes Rendus Mathematique **346(15-16)**, 907–912 (2008)
44. Donaldson, D., Weymark, J.A.: Properties of fixed-population poverty indices. Int. Econ. Rev. **27(3),** 667–688 (1986)
45. Duclos, J.-Y. What is 'pro-poor'? Soc. Choice Welfare **32(1),** 37–58 (2009)
46. Duclos J.-Y., Araar, A.: Poverty and Equity, Measurement, Policy, and Estimation with DAD. Springer, Ottawa (2006)
47. Duclos, J.-Y., Makdissi, P.: Restricted and unrestricted dominance for welfare, inequality and poverty orderings. J. Public Econ. Theory **6**, 145–164 (2004)
48. Duclos, J.-Y., Sahn, D.E., Younger, S.D.: Robust multidimensional spatial poverty comparisons in Ghana, Madagascar, and Uganda. World Bank Econ. Rev. **20(1),** 91–113 (2006)
49. Duclos, J.-Y., Sahn, D.E., Younger, S.D.: Robust multidimensional poverty comparisons. Econ J. **116(514),** 943–968 (2006)
50. Ebert, U., Moyes, P.: A simple axiomatization of the Foster, Greer and Thorbecke poverty orderings. J. Public Econ. Theory **4(4),** 455–473 (2002)
51. Fields, G.S.: Distribution and Development: A New Look at the Developing World. Russell Sage Foundation, New York, and The MIT Press, Cambridge, MA and London (2001)
52. Fields, G.S.: Should poverty and inequality measures be combined? In de Janvry, A., Kanbur, R. (eds.) Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke , pp. 67–74. Kluwer Academic Publishers, Norwell, MA (2005)
53. Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. Working Paper No. 243, Department of Economics, Cornell University (1981)

54. Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. Econometrica **52**, 761–776 (1984)
55. Foster, J.E.: On economic poverty: a survey of aggregate measures. In Basmann, R.L., Rhodes, Jr., G.F.(eds.) Advances in Econometrics, 3, pp. 215–251. JAI Press, Ltd., Greenwich, CT (1984)
56. Foster, J.E.: Absolute versus relative poverty. Amer. Econ. Rev. **88(2),** 335–341 (1998)
57. Foster, J.E.: Poverty indices. In de Janvry, A., Kanbur, R. (eds.) Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke,pp. 41–66. Kluwer Academic Publishers, Norwell, MA (2005)
58. Foster, J.E.: Inequality measurement. In Clark, D. (ed.) The Elgar Companion to Development Studies, pp. 275–280. Edward Elgar, Cheltenham, UK (2006)
59. Foster, J.E.: A class of chronic poverty measures. In Addison, T., Hulme, D., Kanbur, R. (eds.) Poverty Dynamics: Interdisciplinary Perspectives, pp. 59–76. Oxford University Press, Oxford (2009)
60. Foster, J.E., Horowitz, A.W, Mendez, F.: An axiomatic approach to the measurement of corruption: theory and applications. Working Paper 29, OPHI, University of Oxford (2009)
61. Foster, J.E., Sen, A.: On economic inequality after a quarter century. In Sen, A. (ed.) On Economic Inequality, expanded edition, pp. 107–221. Oxford University Press, Oxford (1997)
62. Foster, J.E., Shorrocks, A.F.: Inequality and poverty orderings. Europ. Econ. Rev. **32(2-3),** 654–661(1988)
63. Foster, J.E., Shorrocks, A. F.: Poverty orderings. Econometrica **56(1),** 173–177 (1988)
64. Foster, J.E., Shorrocks, A.F.: Poverty orderings and welfare dominance. Soc. Choice Welfare **5(2-3),** 179–198 (1988)
65. Foster, J.E., Shorrocks, A.F.: Subgroup consistent poverty measures. Econometrica **59(3),** 687–709 (1991)

66. Foster, J.E., Szekely, M.: Is economic growth good for the poor? Tracking low incomes using general means. Int. Econ. Rev. **49**, 1143–1172 (2008)
67. Greer, J.: Food consumption and poverty among small-scale Kenyan farmers. Ph.D. Dissertation, Cornell University, Ithaca, New York (1985)
68. Greer, J., Thorbecke, E.: A methodology for measuring food poverty applied to Kenya. J. Devel. Econ. **24**, 59–74 (1986)
69. Greer, J., Thorbecke, E.: Food poverty profile applied to Kenyan smallholders. Econ. Devel. Cult. Change **35(1),** 115–141 (1986)
70. Greer, J., Thorbecke, E.: Food Poverty and Consumption Patterns in Kenya. International Labor Office, Geneva (1986)
71. Gundersen, C.: Measuring the extent, depth, and severity of food insecurity: an application to American Indians in the USA. J. Population Econ. **21(1),** 191–215 (2008)
72. Haddad, L., Kanbur, R.: How serious is the neglect of intra-household inequality? Econ. J. **100**, 866–881 (1990)
73. Hagenaars, A.: A class of poverty indices. Int. Econ. Rev. **28**, 583–607 (1987)
74. Huppi, M., Ravallion, M.: The sectoral structure of poverty during an adjustment period: evidence for Indonesia in the mid-1980s. World Devel., **19(12),** 1653–1678 (1991)
75. Jalan, J., Ravallion, M.: Is transient poverty different? Evidence for rural china. J. Devel. Stud. **36(6),** 82–99 (2000)
76. Jenkins, S.P., Lambert, P.J.: Three 'i's of poverty curves, with an analysis of UK poverty trends. Oxford Econ. Pap.–New Series **49(3),** 317–327 (1997)
77. Jolliffe, D.: Continuous and robust measures of the overweight epidemic: 1971–2000. Demography **41(2),** 303–314 (2004)
78. Kakwani, N.: Statistical inference in the measurement of poverty. Rev. Econ. Statist. **75(4),** 632–639 (1993)
79. Kakwani, N.: Testing an hypothesis about differences in poverty estimated from grouped data. Econ. Letters **46(3),** 263–271 (1994)

80. Kakwani, N., Silber, J.: Quantitative Approaches to Multidimensional Poverty Measurement. Palgrave Macmillan, New York (2008)
81. Kakwani, N., Silber, J.: The Many Dimensions of Poverty. Palgrave Macmillan, New York (2008)
82. Kanbur, S.M.R.: Measurement and alleviation of poverty, with an application to the effects of macroeconomic adjustment. Int. Monet. Fund Staff Pap. **34(1),** 60–85 (1987)
83. Kanbur, S.M.R.: Transfers, targeting and poverty. Econ. Pol.: A Europ. Forum **4**, 112–47 (1987)
84. Kanbur, S.M.R.: Growth, inequality and poverty: some hard questions. J. Int. Aff. **58(2),** 223–232 (2005)
85. Kanbur, S.M.R.: Poverty and distribution: twenty years ago and now. http://www.arts.cornell.edu/poverty/kanbur/PovAndDistTwentyYears.pdf. Cited 5 Dec 2008 (2008)
86. Kanbur, S.M.R. (ed.): Q-Squared: Qualitative and Quantitative Methods of Poverty Appraisal. Permanent Black, Delhi (2003)
87. Kanbur, S.M.R., Mukherjee, D.: Premature Mortality and Poverty Measurement. Bulletin of Economic Research. **59 (4)**, 339-359 (2007)
88. Kolm, S.C.: The optimal production of social justice. In Margolis, J., Guitton, H. (eds.) Public Economics, pp. 145–200. Macmillan, London. (1968)
89. Kolm, S.C.: Unequal inequalities I. J. Econ. Theory **12**, 416–442 (1976)
90. Kolm, S.C.: Unequal inequalities II. J. Econ. Theory **13**, 82–111 (1976)
91. Kundu, A.: Measurement of poverty—some conceptual issues. Anvesak **11**, 80–96 (1981)
92. Lanjouw, P., Ravallion, M.: Poverty and household size. Econ. J. **105**, 1415–1434 (1995)
93. Levy, S.: Poverty alleviation in Mexico. Working Paper Series #679, World Bank, Washington, DC (1991)
94. Ligon, E., Schechter, L.: Measuring vulnerability. Econ. J. **113(486),** C95–C102 (2003)

95. Lubrano, M., Bauwens, L., Kirman, A., Protopopescu, C.: Ranking economics departments in Europe: a statistical approach. Mimeo (2003)
96. McCulloch, N., Calandrino, M.: Vulnerability and chronic poverty in rural Sichuan. World Devel. **31(3),** 611–628 (2003)
97. Mookherjee, D., and Shorrocks, A.F.: A decomposition analysis of the trend in U.K. income inequality. Econ. J. **92**, 886–902 (1982)
98. Morduch, J.: The microfinance promise. J. Econ. Lit. **37(4),** 1569–1614 (1999)
99. Morduch, J.: The microfinance schism. World Devel. **28(4),** 617–629 (2000)
100. Myles, G. D.: Public Economics. Cambridge University Press, Cambridge (1995)
101. Orshansky, M.: Counting the poor: another look at the poverty profile. Soc. Sec. Bull. **28**, 3–29 (1965)
102. Peichl, A., Schaefer, T., Schleicher, C.: Measuring richness and poverty: a micro data application to Germany and the EU-15. IZA Discussion Paper 3790, October (2008)
103. Ravallion, M.: Poverty Comparisons. Harwood Academic Press, Chur, Switzerland (1994)
104. Ravallion, M., Bidani, B.: How robust is a poverty profile? World Bank Econ. Rev. **8**, 75–102 (1994)
105. Ravallion, M., Chao, K.: Targeted policies for poverty alleviation under imperfect information - algorithms and applications. J. Pol. Modeling **11(2),** 213–224 (1989)
106. Ravallion, M., Chen, S.: Weakly relative poverty. World Bank Policy Research Working Paper 4844, World Bank, Washington, DC (2009)
107. Ray, A.: A class of decomposable poverty measures with public transfers. Econ. Bull. **9**, 1–8 (2006)
108. Ray, D.: Development Economics. Princeton University Press, Princeton, NJ (1998)
109. Rodgers, J.L.; Rodgers, J.R.: Measuring the intensity of poverty among subpopulations: applications to the United States. J. Human Res. **26(2),** 338–361 (1991)

110. Rodgers, J.R., Rodgers, J.L.: Chronic poverty in the United States. J. Human Res. **28(1),** 25–54 (1993)

111. Rowntree, B.S.: Poverty: A Study of Town Life. Macmillan, London (1901)

112. Runciman, W.G.: Relative Deprivation and Social Justice. Routledge, London (1966)

113. Sahn, D.E., Stifel, D.C.: Robust comparisons of malnutrition in developing countries. Amer. J. of Agr. Econ. **84(3)**, 716–735 (2002)

114. Seidl, C.: Poverty measurement: A survey. In Bos, D., Rose, M., Seidl, C. (eds.) Welfare and Efficiency in Public Economics, pp. 71–147. Springer-Verlag, Berlin/New York (1988)

115. Sen, A.: Poor relatively speaking. Oxford Econ. Pap. **35**, 153–169 (1983)

116. Sen, A.: Conceptualizing and measuring poverty. In Grusky, D.B., Kanbur, S.M.R. (eds). Poverty and Inequality, pp. 30–46. Stanford University Press, Palo Alto, CA (2006)

117. Sen, A.K.: On Economic Inequality. Clarendon Press, Oxford (1973)

118. Sen, A.K.: Poverty: an ordinal approach to measurement. Econometrica **44**, 219–231 (1976)

119. Sen, A.K.: Issues in the measurement of poverty. Scand. J. Econ. **81**, 285–307 (1979)

120. Shorrocks, A.F.: The class of additively decomposable inequality measures. Econometrica **48**, 613–625 (1980)

121. Shorrocks, A. F.: Aggregation issues in inequality measurement. In W. Eichhorn (ed.) Measurement in Economics, pp. 429–451. Physica-Verlag, Heidelberg (1988)

122. Shorrocks, A.F.: Revisiting the Sen Poverty Index. Econometrica **63**, 1225–1230 (1995).

123. Son, H.H., Kakwani, N.: Poverty equivalent growth rate. Rev. Income Wealth **54**, 643–655 (2008)

124. Subramanian, S.: Measurement of Inequality and Poverty, Oxford University Press, Delhi (1997)

125. Subramanian, S.: A Re-scaled version of the Foster-Greer-Thorbecke Poverty Indices based on an association with the Minkowski Distance Function, UNU-WIDER Research Paper RP2004/10 (2004)
126. Tarp, F., Arndt, C., Jensen, H. T., Robinson, S., Heltberg, R.: Facing the development challenge in Mozambique: an economywide perspective. Research reports 126, International Food Policy Research Institute (IFPRI), Washington, DC (2002)
127. Theil, H.: Economics and Information Theory. North-Holland, Amsterdam (1967)
128. Thon, D.: On measuring poverty. Rev. Income Wealth **25**, 429–440 (1979)
129. Thorbecke, E.: Conceptual and measurement issues in poverty analysis.UNU/WIDER Working Paper Series (2004)
130. Thorbecke, E.: Multidimensional poverty: conceptual and measurement issues. In Kakwani, N., Silber, J. (eds.) The Many Dimensions of Poverty, pp. 3–19. Palgrave Macmillan, New York (2008)
131. Thorbecke, E., Berrian, D.: Budgetary rules to minimize societal poverty in a general equilibrium context. J. Devel. Econ. **39(2)** 189–205 (1992)
132. Thorbecke, E., Jung, H.-S.: A multiplier decomposition method to analyze poverty alleviation. J. Devel. Econ. **48**, 279–300 (1996)
133. Todaro, M.P., Smith, S.C.: Economic Development, 10th ed. Addison-Wesley/Pearson, Boston (2008)
134. Tsui, K.Y.: Multidimensional poverty indices. Soc. Choice Welfare **19(1),** 69–93 (2002)
135. Watts, H.W.: An economic definition of poverty. In Moynihan, D. P. (ed.) On Understanding Poverty, pp. 316–329. Basic Books, New York (1968)
136. Xu, K.: U-statistics and their asymptotic results for some inequality and poverty measures. Econometric Rev. **26(5),** 567–577 (2007)
137. Zheng, B.: Can a poverty index be both relative and absolute. Econometrica **62(6),** 1453–1458 (1994)
138. Zheng, B.: Aggregate poverty measures. J. Econ. Surveys **11**, 123–162 (1997)

139. Zheng, B.: Poverty orderings. J. Econ. Surveys **14(4),** 427–466 (2000)

140. Zheng, B.: Statistical inference for poverty measures with relative poverty lines. J. Econometrics **101**, 337–356 (2001)

141. Zheng, B.: Poverty comparisons with dependent samples. J. Appl. Econometrics **19(3),** 419–428 (2004)

FIGURE CAPTION

Figure 1. P_2 and the distribution