POVERTY, PRICES, AND PLACE: HOW SENSITIVE IS THE SPATIAL DISTRIBUTION OF POVERTY TO COST OF LIVING ADJUSTMENTS?

DEAN JOLLIFFE*

This article examines how accounting for cost-of-living differences across metropolitan and nonmetropolitan areas affects measured rates of poverty. The spatial price index used is based on the Fair Market Rent data and was developed by the Census Bureau for use in its experimental poverty research program. Following U.S. federal definitions, poverty in nonmetro areas has been consistently higher than it has been in metro areas. Using the Fair Market Rent index to adjust for differences in cost of living results in a complete reversal of nonmetro-metro rankings in terms of prevalence, depth, and severity of poverty for every year examined (1991 to 2002). (*JEL* I32, R1, C81)

I. INTRODUCTION

Estimates of poverty figure prominently in the criteria that determine the geographic distribution of large sums of cash and in-kind benefits from state and federal government programs. Citro and Michael (1995, 89–90) note that in the early 1990s, 27 different federal assistance programs linked their eligibility criteria in part to poverty lines or area average poverty rates. As one example, an eligibility criterion for the Food Stamp Program is that household income must be equal to or less than 130% of the poverty line. In 2003, the Food Stamp Program distributed $21 billion in program benefits, and data from the 2003 Current Population Survey indicate that per capita benefits were 39% higher in nonmetropolitan areas than they were in metropolitan areas of the United States.[1] As another example, federal block grants for community development are typically linked to county-level poverty estimates; Reeder (1996, 1) notes that persistently poor nonmetro counties benefit disproportionately from block grants, receiving more than $1,000 per person in 1994. Reeder et al. (2001, 4) show that in 1997 the per capita distribution of federal funds for social safety net programs was 8% higher in nonmetro areas than in metro areas.

This distribution of social welfare assistance is arguably well targeted because the estimated prevalence of poverty has been greater in nonmetro areas than in metro areas in every year since the 1960s, when poverty rates were first recorded; see Jolliffe (2003b). However, this argument is potentially sensitive to the following issues: first, how poverty is defined plays an important role in the geographic

*I wish to thank Kathleen Short and Erika Steinmetz of the Census Bureau for their help in providing me with the FMR index for this analysis. I thank Alfred Meier of the Census Bureau for answering many questions regarding CPS sample design and interpretation of Census Bureau publications. I also thank for comments: Constance Newman, Laura Tiehen, Leslie Whitener, and session participants at the 2004 American Social Sciences Association meetings and the 2004 Western Economic Association International conference. The views and opinions expressed in this article do not necessarily reflect the views of the Economic Research Service of the U.S. Department of Agriculture. I assume responsibility for the contents of this article and any errors it may contain.

Jolliffe: Economic Research Service, Room N-2188, 1800 M Street NW, Washington, DC 20036. Phone 1-202-694-5430, Fax 1-202-694-5642, E-mail Jolliffe@ers.usda.gov

1. Per capita benefits are averaged over the entire population (recipients and nonrecipients). This finding is similar to that of Ghelfi (2003), who uses data from the Bureaus of Economic Analysis and documents that per capita food stamp benefits were 32% greater in nonmetro areas than metro areas during 2001.

ABBREVIATIONS

CPS: Current Population Survey
FGT: Foster-Greer-Thorbecke
FMR: Fair Market Rent
NAS: National Academy of Sciences
PSU: Primary Sampling Unit

Economic Inquiry
(ISSN 0095-2583)
Vol. 44, No. 2, April 2006, 296–310

doi:10.1093/ei/cbj016
Advance Access publication November 16, 2005

distribution of measured poverty and, second, any changes to the definition could affect the geographic distribution of poverty and funding for social safety nets.

The National Academy of Sciences (NAS) Panel on Poverty and Family Assistance has recommended several changes in how the federal government measures poverty, including adjustments for geographic differences in the cost of living—see Citro and Michael (1995, 186–201). Although the NAS panel recommends several changes, Nord and Cook (1995) show that adjusting for cost-of-living differences is the one aspect of reform that would most systematically change the geographic distribution of poverty. Currently, the official federal poverty thresholds assume that the cost of living is the same over the entire United States, but the Census Bureau has developed experimental poverty measures that use Fair Market Rent (FMR) data to create an index for spatial differences in the cost of living; see Short (2001a, A2–A5). The purpose of Census Bureau's research into experimental poverty measures is to present its attempts at implementing the NAS recommendations with the aim of improving the official measure of poverty; see Short (2001b, 1). The purpose of this article is to examine how the use of the FMR index currently under consideration by the Census Bureau affects the relative distribution of poverty across metro and nonmetro areas of the United States.

The focus on metro and nonmetro areas is driven largely by the fact that an adjustment for cost-of-living differences will have the greatest effect on this comparison and also because of the strong historical difference in poverty rates across these areas. Furthermore, this geographic distinction mirrors the design of the FMR index, and the disbursement of federal funds is in some cases linked to whether a county is identified as being metro or nonmetro.

Because this article examines solely how spatial price differences affect poverty rates, it provides information on only one of the NAS panel's suggested changes, and it is important to interpret the results with this caveat in mind. Nonetheless, an advantage to this narrow focus on spatial-price adjustments is that the findings readily highlight the sensitivity of the relative poverty levels of nonmetro and metro areas to this change. The results from this analysis suggest a complete reversal of all poverty measures considered. Specifically, once adjusted for cost-of-living differences using the FMR index, metro poverty is greater than nonmetro poverty in terms of prevalence, depth, and severity over the entire period considered in this analysis (1991–2002).

This article adds to the current literature on the reform of the poverty measure in three ways. First, it focuses on the impact of change on relative differences in poverty rates between metro and nonmetro areas. This focus is important both in terms of understanding how reform could affect the geographic distribution of benefits from federal assistance programs and also in terms of the potential political economy issues that might develop from such a proposed change.

Second, much of the current analysis of the experimental poverty measures is based on how change will affect the prevalence of poverty. In this article, I consider four measures of poverty that help to establish the robustness of the findings as well as shed light on the distributional effects of the proposed change. Three of the measures belong to the Foster-Greer-Thorbecke (FGT) family of poverty measures and have been widely used in the international poverty literature—see Foster et al. (1984).[2] The headcount, or P_0, is the standard measure used and provides a measure of the prevalence of poverty. The poverty gap, or P_1, measure provides a measure of the depth of poverty; and the severity of poverty index, or P_2, is sensitive to changes in the income distribution of the poor. The fourth measure is the Watts index. For examples and a discussion of the importance of considering poverty indices other than just the headcount index, see Zheng's meticulous survey on poverty measures (1997).

The final way in which this article adds to the experimental poverty literature is in terms of methodology. The statistical tests for nonmetro-metro poverty differences are corrected for features of the sample design.[3]

2. The following all use these three measures: Jolliffe et al. (2004) detail Egypt; Datt and Ravallion (1992) cover Brazil and India; Howes and Lanjouw (1998) use examples from Pakistan and Ghana; Kakwani (1993) examines Côte d'Ivoire; and Ravallion and Bidani (1994) examine Indonesia.

3. Zheng (2001) provides design-corrected estimates of sampling variance for poverty estimates based on relative poverty lines (i.e., the poverty line is relative to the distribution of income, such as one-half the median income level). The advantage of the estimates provided in this article is that they are based on a fixed (or absolute) poverty line, which is how poverty is measured in the United States. Another advantage is that Jolliffe and Semykina (1999) provide a Stata program that estimates the standard errors presented in this article.

Most nationally representative data sets, particularly those from which poverty estimates are formed, are not based on simple random draws from the population; rather, they are frequently based on complex sample designs. As one example, the sample used for the Current Population Survey (CPS) is drawn from a census frame using a stratified, multistage design. Howes and Lanjouw (1998) present evidence that estimated standard errors for poverty measures can have large biases when false assumptions are made about the nature of the sample design. In particular, they show that if the sample design is multistaged but that standard errors are derived from the incorrect assumption of a simple random sample, then the standard errors will significantly underestimate the true sampling variance. An example from Jolliffe et al. (2004) show that in the case of poverty measures for Egypt, failing to adjust for the characteristics of the sample design would result in an underestimate of the correct standard errors by 187% to 212%.

The remaining part of this article proceeds as follows. Section II covers poverty measurement issues and includes a discussion of the Fair Market Rents index, the data, poverty line, poverty measures, and the estimates of sampling variance.[4] Section III provides a discussion of the results. The primary finding is that without correcting for spatial-price differences, nonmetro poverty is higher than metro poverty. Using the FMR index reverses this result during all 12 years examined. Section IV provides a brief conclusion.

II. POVERTY MEASUREMENT

The FMR Index

Although the data are limited, the evidence suggests that there are significant geographic differences in the cost of living. Up until 1982, the Bureau of Labor Statistics collected data from the Family Budget Program, which provided estimates of the relative cost of a consumption bundle for a family of four living in different areas of the United States. The last sample of the Family Budget Program data indicates that in 1981 there was significant spatial variation in the cost of purchasing the fixed bundle of goods. For example, in urban areas of Northwest the cost was 113% of the national average, and in the nonmetro South it was 91% of the national average—see Citro and Michael (1995, 186).

Currently, there are other sources of data on spatial differences in the cost of living, but none of these include both metro and nonmetro regions of the United States. For example, the Bureau of Labor Statistics has been working on using data from the consumer price index to develop a spatial price index, but these efforts have focused strictly on metropolitan areas; note the work of Kokoski (1991) and Moulton (1995) on this point. Perhaps the best-known spatial price index is the ACCRA index, developed by the American Chamber of Commerce Researchers Association.[5] The primary shortcoming of this index is that it provides an estimate of only the cost of living in an area if a volunteer has reported data and that it is intended to measure differences among urban areas. Further, the ACCRA index is designed to reflect cost-of-living differences for a professional household and is based on the typical spending patterns of households in the top quintile of income. For the purposes of poverty analysis, this is not a useful index.

This article uses the spatial price index currently under consideration by the Census Bureau in its research on experimental poverty measures—see Short (2001a, 2001b)—and is based on the FMR data collected by the U.S. Department of Housing and Urban Development.[6] The primary advantages of these data are that they provide full coverage of the United States and reflect spending of lower-income households. Housing and Urban Development produces annual estimates of the FMR for 354 metro areas and 2,350 nonmetro counties. The FMR data estimate the cost of gross rent (utilities included) at the 40th percentile for "standard" quality housing.[7] The purpose of the FMR is to determine eligibility of rental housing units for the Section 8

4. Parts of this section are drawn from Jolliffe (2003b).

5. For an ACCRA example, see Dumond et al. (1999). See Koo et al. (2000) for a comparison of a consumer price index and the ACCRA index.

6. For a critique of the FMR as a spatial price index, see Short (2001b, appendix A). For an alternate examination of spatial price differences between metro and nonmetro areas, see Nord (2000).

7. From 1995 to 2001, the FMR is based solely on the 40th percentile. As of 2001 the FMR index is based on the 40th percentile except for 39 metropolitan statistical areas which are based on the 50th percentile. Between 1983 and 1994, the index was based on the 45th percentile.

Housing Assistance Payments program. Section 8 participants cannot rent a unit if the rent exceeds the FMR. (FMR also serves as the payment standard used to calculate subsidies under the Rental Voucher program.) See U.S. Housing and Urban Development (2005) for more details.

The FMR index used in this analysis, constructed by Short (2001a, 2001b) for the Census Bureau's experimental poverty measures, is a fixed-weight index consisting of two components—housing and all other goods and services. The index, following the recommended approach of the National Academy of Sciences report—see Citro and Michael (1995, 197–98)—assigns a weight of 44% for housing expenses and 56% for all other goods and services. Further, the index assumes that variation in the FMR data reflects variation in housing prices for the poor and that there is no variation in prices of all other goods and services. The focus on housing prices in the index is supported by Moulton (1995, 181), who notes that "the cost of shelter is the single most important component of interarea differences in the cost-of-living." By construction then, if the FMR data indicate that rents in a particular area are 10% higher than the baseline, then the FMR index used by the Census Bureau reflects a cost of living in this area that is 4.4% higher than the baseline.

The justification for assuming no spatial variation in prices for the nonhousing goods and services in the consumption bundle of the poor is based primarily on two reasons. First, the lack of credible data sources for spatial price variation in nonrent prices is an important factor. Second, a common presumption is that high-rent areas are somewhat higher in the prices of other goods (or that the correlation in prices is positive) so that this index will present a conservative, or lower bound, estimate of the true variation in prices. Section III examines the sensitivity of the core findings to the assumption of no correlation.

Finally, the FMR index used in the Census Bureau's experimental poverty measures is aggregated up to 100 different price levels, one for metro and one for nonmetro areas of each state plus the District of Columbia.[8] I have also scaled the FMR index to ensure that the FMR-adjusted poverty estimates match the official U.S. federal estimates at the national level. With this scaling, any deviation from official estimates at the subnational level will be strictly due to relative price differences in the index. Comparing area averages of this scaled index suggests that the cost of living in nonmetro areas is 79% that of metro areas. This is the first indication that using this index to adjust for cost-of-living differences is likely to have a significant effect in terms of measuring metro-nonmetro poverty differences.

The Data: 1992–2003 CPS and the U.S. Poverty Thresholds

The data used in this article are from the 1992 through 2003 March supplement to the CPS, which is conducted by the Bureau of the Census for the Bureau of Labor Statistics. The CPS data form the basis for the official U.S. poverty estimates and provide information on approximately 50,000 households in each year. The March supplement, also called the Annual Demographic Survey of the CPS, collects information on income and a variety of demographic characteristics. The reference period for income-related questions is the preceding calendar year; therefore, the 1992 to 2003 CPS data provide poverty estimates for 1991 through 2002.

The sample is representative of the civilian, noninstitutionalized population and members of the armed forces either living off base or with their families on base. The sample frame is based on housing structures, not individuals; so all individuals who are homeless at the time of the interview are excluded from the sample. Because the homeless are disproportionately located in metro areas, their exclusion disproportionately biases the metro poverty estimates downward. A primary finding of this article is that adjusting for cost-of-living differences with the FMR index decreases nonmetro and increases metro poverty to the extent that the nonmetro-metro poverty rankings are completely reversed. If the homeless were included in this analysis, they would further reinforce this reranking of poverty.

The measure of welfare used in this article is income as it defined for federal poverty rates. This definition includes all pretax income but does not include capital gains or any noncash benefits, such as public housing, medicaid, or food stamps. The poverty thresholds used in

8. New Jersey and the District of Columbia consist of only metro areas—hence, 100 total FMR price levels.

this article are the U.S. federal government poverty lines, which were developed in 1965 following a cost-of-basic-needs methodology that sets the poverty line at the value of a consumption bundle considered to be adequate for basic consumption needs. Basic needs, in this context, represent a socially determined, normative minimum for avoiding poverty. For more details on this methodology and other methods of drawing poverty lines, see Ravallion (1998).

The U.S. poverty line set in 1965 was based on the cost of the U.S. Department of Agriculture's economy food plan, a low-cost diet determined to be nutritionally adequate. In addition to the cost of this food plan, the poverty line included an allowance for nonfood expenditures that was twice the value of the cost of the department's economy food plan.[9] To account for inflation, the poverty lines set in 1965 are adjusted each year using a price index.[10] The latest poverty line used in this study is from 2002, and it is set at \$9,359 for an individual under 65 years of age, \$12,400 for a two-person family with one child and one adult, and \$21,469 for a family with two adults and three children. For a complete listing of 2002 poverty lines for individuals and families of various sizes, see Proctor and Dalaker (2003, 4).

The Poverty Measures

The previous section describes the measure of welfare and poverty lines used to identify who is poor. The next step is to aggregate this information into a scalar measure of poverty. To examine the sensitivity of estimated poverty levels to the choice of a poverty measure, I consider three measures that belong to the FGT family and also the Watts measure. The first FGT measure is the headcount index (P_0), which is the percentage of the population that is poor. The second measure, called the poverty gap index (P_1), is found by first measuring the income gap (i.e., the proportionate difference between income and the poverty line) for all poor persons. The poverty gap index is then the average value of the income gaps, where the average is formed over the entire population, counting the nonpoor as having zero income gap. The third measure is P_2, which is said to reflect the severity of poverty and is defined as the mean value of the squared income gaps.

The FGT class of poverty measures, also referred to as P_α, can be represented as

$$(1) \qquad P_\alpha = 1/n \sum_i I(y_i < z)[(z - y_i)/z]^\alpha$$

where n is the sample size, i subscripts the family or individual, y is the relevant measure of welfare, z is the poverty line, and I is an indicator function that takes a value of 1 if the statement is true and 0 otherwise. When $\alpha = 0$, the resulting measure is the headcount measure, or P_0. When $\alpha = 1$, the FGT measure results in the poverty gap measure, or P_1, and P_2 results when $a = 2$.

The usefulness of these measures can be illustrated by considering a transfer of money from a rich person to a poor person that is not large enough to push the poor person over the poverty line. This transfer has no effect on the headcount measure, P_0, but the poor person is better off and this welfare improvement is reflected in a reduction of both P_1 and P_2. As another example, a transfer of income from a poor person to a poorer person will not alter either the headcount or the poverty gap measure, but it improves the distribution of income of the poor; this change is reflected by a reduction in P_2.[11]

These examples point to an important reason to consider the poverty gap (P_1) and the P_2 measures in addition to the commonly reported headcount measure (P_0). A frequently stated goal of many programs is the reduction of poverty, but the policies that are appropriate to attain this goal will vary depending on which poverty measure is considered. If policymakers are focused on the headcount measure, then the most efficient way to reduce poverty is through assistance

9. For details on the first poverty lines, see Orshansky (1965). For a history of poverty lines before Orshansky, see Fisher (1997). For a critical discussion of the official poverty line, see Ruggles (1990).

10. Before 1969, the index used was the changing cost of the U.S. Department of Agriculture economy food plan; afterward, the consumer price index for all goods and services has been used.

11. Unlike the Sen (1976) or Kakwani (1980) distribution-sensitive measures of poverty, the P_2 measure satisfies the "subgroup consistency" property, which means that if poverty increases in any subgroup and does not decrease elsewhere, then aggregate poverty must also increase—see Foster and Shorrocks (1991).

to the least poor. If, on the other hand, policy-makers are concerned about the overall welfare of the poor and not just on reducing the number of persons living in poverty, then the appropriate measure is one that captures the depth (P_1) and severity (P_2) of poverty.

While the P_2 poverty measure is distribution sensitive, it does not satisfy the more demanding poverty axiom of "weak transfer sensitivity."[12] To satisfy this axiom, a poverty measure must be more sensitive to transfers between poor and poorer persons at the lower end of the income distribution. Consider two transfers of a dollar, both from a poor to a poorer person, and in both cases the difference of incomes between the two persons is the same. Assume that in the first case, the transfer is occurring between two people who are in more severe poverty (further from the poverty line) than are those in the second case. A poverty measure respecting weak transfer sensitivity would fall more for the first case. A measure satisfying this axiom, and the one considered in this article, is the Watts measure, which can be expressed as

$$(2) \qquad W = 1/n \sum_i I(y_i < z)[\ln(z) - \ln(y_i)]$$

where the terms are defined as they are in equation 1.[13]

Estimating Standard Errors for the Poverty Measures Based on CPS Data

To test whether adjusting for cost-of-living differences affects the relative poverty rankings requires estimates of the sampling variance for the measures. Kakwani (1993) provides asymptotic variance estimates for the FGT povertymeasures that are easy to calculate and are frequently used. The Kakwani formula for the variance of P_0, the headcount measure, is $P_0(1 - P_0)/(n - 1)$, where n is the sample size. The formula for all other variance estimates of the FGT measures is$(P_{2\alpha} - P_\alpha^2)/(n - 1)$. The primary disadvantage of the Kakwani estimates is that they assume that the sample is a simple random sample of the population.

As noted in the introduction, using the Kakwani standard errors when the data were collected using a multistage sample design can result in a large underestimate of the true sampling variance. The strategy used in this article to estimate the sampling variance corrected for design effects is to first derive exact (analytical) estimates for the poverty measures and then to address the issue of sample design. An advantage of the Watts and FGT poverty measures in this context is that they are additively decomposable, a characteristic that greatly simplifies deriving the analytical estimates of the sampling variance of the poverty measures. To illustrate, consider any income vector y, broken down into m subgroup income vectors, $y(^1), \ldots, y(^m)$. Because P_α is additively decomposable with population share weights, it can be written as

$$(3) \qquad P_\alpha(y; z) = \sum_{j=1}^{m} (n_j/n) P_{\alpha,j}(y^j; z)$$

where n is the sample size, n_j is the size of each subgroup, and z is again the poverty line. By treating each observation as a subgroup, the estimate of poverty is the weighted mean of the individual-specific measures of poverty, and the sampling variance of the poverty measure is the variance of this mean, or

$$(4) \qquad P_\alpha = \sum_{i=1}^{n} P_{\alpha,i}/n \quad \text{and}$$

$$V(P_\alpha) = n^{-1}(n-1)^{-1} \sum_{i=1}^{n} (P_{\alpha,i} - P_\alpha)^2$$

where i is the individual.

The next step is to incorporate the sample design information, which typically requires that the researcher have access to not only unit record data but also data identifying the characteristics of the sample design. In the case of the CPS data, the sample design information that identifies the strata and primary sampling units (PSUs) has been censored from the public-use files to maintain respondent confidentiality. To compensate for the missing design information, the U.S. Census Bureau (2000, appendix C) provides detailed notes on how to approximate design-corrected

12. Zheng (2000, 120–23) notes that the FGT indices have the somewhat unattractive feature of exhibiting increasing poverty aversion in income (for income below the poverty line). In other words, the FGT indices become less distribution sensitive when evaluated at lower levels of income for all levels of income below the poverty line.

13. An important disadvantage of the Watts measure for this analysis is that it is undefined for nonpositive income levels. The CPS data contain observations with zero recorded income and comprise about 1% of the sample in each of the years.

standard errors for a limited set of poverty estimates. An important shortcoming of this method is that parameter estimates are provided for only the headcount measure; there are no corrections provided for any other measures of poverty.[14]

In addition to the issue that the Census Bureau does not provide sample-design corrections for the P_1, P_2, or Watts poverty measures, an additional problem is that the recommended method appears to be significantly less precise for nonmetro-metro comparisons. The proposed correction for all nonmetro statistics provided by U.S. Census Bureau (2000, appendix C) is to multiply the design-correction coefficients by 1.5. The implication of this correction is that for all statistics the ratio of the design effects for metro to nonmetro areas is constant. Another factor likely to affect the accuracy of this correction is that it has not been updated in the last 20 years, whereas the design-correction coefficients for all other characteristics are frequently updated.[15]

Given that the Census Bureau–recommended method does not provide corrections for the sampling variance of P_1 and P_2 and that the adjustment factor for nonmetro areas appears to be a rough approximation, I abandon this method. Instead, I follow an approach based on replicating aspects of the CPS sample design by creating synthetic variables for the strata and clusters that induce similar design effects. A more detailed description of the approach and simulation results suggesting that it provides useful approximations are provided in Jolliffe (2003a).

14. Another shortcoming of the Census Bureau–recommended method is that corrections are provided for only a limited set of characteristics. For example, the U.S. Census Bureau (2000, appendix C) provides parameter estimates to adjust the sampling variance for the headcount measure by several age categories. If the analysis is focused on individuals 15 to 24 years old, the analyst is provided with parameter estimates. If the relevant subsample is, say, working-age adults, then the Census Bureau does not provide the necessary parameters to estimate standard errors.

15. Personal communication with the Census Bureau appears to support this assertion that the nonmetro adjustment is less precise: "The factor of 1.5 has been used for nonmetro areas as a simple approximation. While the best factor likely varies from characteristic to characteristic, we use 1.5 for all characteristics rather than publishing a different factor for each estimate. Years ago, someone looked at the data for metro/nonmetro areas and decided that 1.5 would be a good, and somewhat conservative, estimate for most characteristics."

The first step of the synthetic design approach for this analysis of poverty is to sort the data by income,[16] then assign each set of four consecutive housing units to a separate cluster. The purpose of the sorting is to induce a high level of intracluster correlation, and the choice of four units matches, on average, the actual CPS cluster size. I select the four regions of the United States as synthetic strata to capture the geographic aspect of the CPS stratification. Appendix Table A-1 provides a summary table from Jolliffe (2003a) illustrating that the synthetic design approach matches the estimates provided by the Census Bureau for the headcount measure.

With the selection of the synthetic strata and clusters, one can then directly obtain design-corrected estimates of sampling variance based on equation 4. Following Kish (1965) and noting from earlier that P_α can be considered a sample mean, the estimated sampling variance of the FGT poverty measures from a weighted, stratified, clustered sample is given by

$$(5)\quad V(P_{\alpha,w}) = \sum_{h=1}^{L} n_h(n_h-1)^{-1} \sum_{i=1}^{n_h} \times \left(\sum_{j=1}^{m_{h,i}} w_{h,i,j} P_{\alpha,h,i,j} - \sum_{i=1}^{n_h} \sum_{j=1}^{m_{h,i}} w_{h,i,j} P_{\alpha,h,i,j} \right)^2$$

where the h subscripts each of the L strata, i subscripts the cluster or PSU in each stratum, j subscripts the ultimate sampling unit, so w_{hij} denotes the weight for element j in PSU i and stratum h. The number of PSUs in stratum h is denoted by n_h, and the number of ultimate sampling units in PSU (h, i) is denoted by m_{hi}.[17]

III. RESULTS

Nonmetro-Metro Comparisons of FMR-Adjusted Poverty Estimates

Before answering whether nonmetro-metro poverty comparisons are sensitive to spatial

16. The methodology requires sorting the data on the variable most relevant to the analysis.

17. The poverty and sampling variance estimates are documented in more detail in Jolliffe and Semykina (1999), which provides a program to estimate equations 1 and 5 in the Stata software.

price adjustments, I need to establish the baseline poverty measures for comparison. Between 1991 and 2002, the nonmetro headcount measure ranges from a high of 0.17 in 1993 (representing 9.7 million poor people) to a low of 0.13 in 2000 (6.8 million people). The metro headcount measure ranges from a high of 0.15 in 1993 (29.5 million people) to a low of 0.11 in 2000 (24.3 million people) living in poverty.

The variation in P_1 and P_2 is similar. Across both these measures, for metro and nonmetro areas alike, poverty was at its lowest level in 2000. In terms of P_1, the year with the highest level of poverty came in 1993. The worst year, as measured by P_2, came in 1997 for nonmetro areas and 1993 for metro areas. In every year and for each measure, nonmetro poverty was worse than the poverty for metro areas.[18] Also in terms of correcting for the complex sample design, over the 72 FGT poverty measures estimated (P_0, P_1, and P_2 for each year from 1991 to 2002 by metro and nonmetro areas), in no case was the design effect less than 4. The implication is that the design-corrected standard errors are all more than twice as large as those that would be estimated if one (incorrectly) ignored the complex sample design.[19]

One interpretation of the poverty gap measure is that it is equal to the product of the headcount measure and the income gap, where the income gap is the average shortfall of the poor as a fraction of the poverty line. This implies that in 1990 the average shortfall of the poor as a fraction of the poverty line is equal to 40% in nonmetro areas and 44% in metro areas. In 2002, the average shortfall in nonmetro areas is equal to 44% of the poverty line, whereas in metro areas this shortfall is 47%. During all 12 years, the average shortfall is greater in metro areas than in nonmetro areas, which indicates that on average the metro poor are worse off than the nonmetro poor.

This difference in the average income shortfall of the poor suggests that there could be differences in the well-being of the poor across areas. Figure 1 explores this issue by graphing density estimates of the welfare ratio (sometimes called the *income-to-needs ratio*, which is the ratio of income to the poverty threshold). The advantage of welfare ratios over income is that they provide measures of well-being that control for demographic differences (and these demographic characteristics may differ across areas).[20] The reason is that they are a function of the poverty thresholds, which are adjusted to reflect different levels of need for families of various size and age.

Figure 1 provides kernal density estimates of metro and nonmetro welfare ratios for 1992, 1995, 1998, and 2001. For all years, the nonmetro welfare ratio is more peaked near the poverty line, indicating that a larger proportion of the nonmetro poor subsist on greater welfare ratios and are therefore relatively better off. Similarly, the nonmetro welfare ratio lies below the metro distribution on the left side of the distribution, indicating that a larger proportion of the metro poor live in extreme poverty. One implication is that a small increase in income would disproportionately move more nonmetro poor persons than metro persons over the poverty line.

Table 1 provides the first comparison of the baseline poverty measures for 2001 and 2002 with poverty measures that have been adjusted for cost-of-living differences using the FMR index. The headcount measure for 2001 shows that 14.2% of nonmetro residents are poor, compared with 11.1% of the metro population. In other words, the nonmetro prevalence of poverty is 28% greater than the metro rate. This ranking holds for the other poverty measures. The 2001 poverty gap is 21% greater in nonmetro than in metro areas; P_2 is 18% higher; and the Watts measure is 34% higher in nonmetro areas. Table 1 also shows that this pattern continues in 2002. Nonmetro poverty is higher than metro poverty across all four measures, though the percentage difference

18. A table of these results is available from the author upon request.

19. The largest design effect is 6.1 for the 2001 nonmetro P_1 measure. This means that the corrected standard errors are almost two and a half times greater than what one would estimate if ignoring the sample design.

20. For example, in 1999 the average age of a metro poor person was 28 years, compared to 32 years for the nonmetro poor. In terms of family size, 16% of the metro poor live in two-person families, compared to 20% for the nonmetro poor.

FIGURE 1
Welfare Ratios of the Poor from 1992 to 2001 by Area

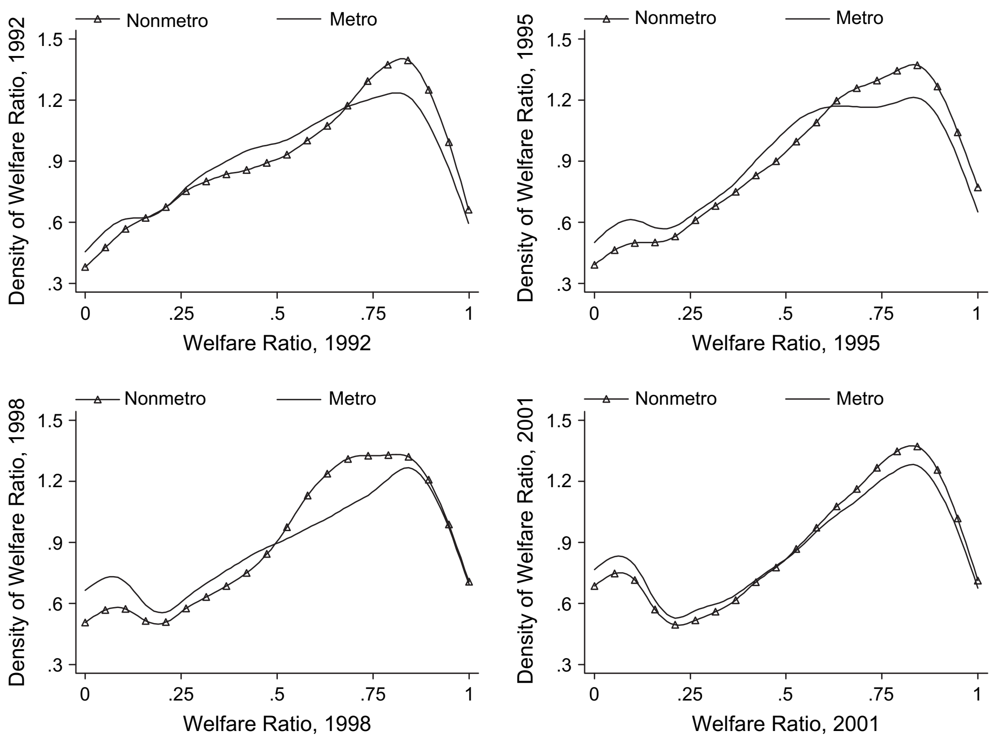

Note: Kernal density estimates of metro and nonmetro welfare ratios (income divided by the poverty line) are for 1992 (upper left panel), 1995 (upper right panel), 1998 (lower left panel), and 2001 (lower right panel). The nonmetro density estimate is marked with triangles. The density of the welfare ratio is measured in terms of the reciprocal of the welfare ratio (not measured on a probability scale) and thus can exceed 1. The Epanechnikov kernal is used for all estimates with a smoothing parameter set to 0.08.

declines as one considers the distribution-sensitive measures, P_1 and P_2.

The estimates listed in the FMR-adjusted columns are the poverty measures when each are calculated based on income levels that have been corrected for spatial differences following the FMR index. In 2001, the official nonmetro poverty rate of 14.2% drops significantly to 10.5% when corrected for spatial-price differences. At the same time, the metro poverty rate increases from 11.1% to 12.0% when adjusted following the FMR index. The net effect is that the prevalence of nonmetro poverty is 12% lower than the metro poverty rate when both measures are adjusted for cost-of-living differences (as measured by the FMR index). Table 1 indicates that this reversal of the relative ranking of nonmetro and metro poverty also holds for the P_1, P_2, and Watts measures in 2001 and 2002.[21]

To understand whether this rather striking reversal of the poverty rankings is unique to recent events, I repeat this analysis for all years between 1991 and 2002.[22] Panel A of Figure 2

21. The analysis focuses on metro and nonmetro areas in part because doing so reflects the design of the FMR index and because several government assistance programs are linked to this geographic definition. Cushing and Zheng (2000) and Jolliffe (2003b) examine the geographic distribution of several poverty measures using central city, suburb, and nonmetro as the geographic units and find that relative poverty rates tend to be the highest in central cities and lowest in suburbs. Adjusting these measures with the FMR index does not result in a reranking of this ordering—FMR-adjusted poverty is highest in central cities and lowest in suburbs.

22. A table of these results is available from the author upon request.

TABLE 1
Nonmetro-Metro Poverty Comparisons, 2001 and 2002

	P_0 Headcount		P_1 Poverty Gap		P_2 FGT ($\alpha = 2$)		Watts Measure	
	Actual	FMR Adjusted	Actual	FMR Adjusted	Actual	FMR Adjusted	Actual	FMR Adjusted
2001								
Nonmetro	0.142	0.105	0.063	0.050	0.043	0.036	0.101	0.079
	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)
Metro	0.111	0.120	0.052	0.055	0.036	0.038	0.075	0.800
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Nonmetro-metro difference	28%	−12%	21%	−9%	18%	−5%	34%	−1%
	(3.80)	(3.01)	(4.71)	(3.97)	(5.64)	(4.97)	(8.20)	(6.91)
2002								
Nonmetro	0.142	0.105	0.062	0.049	0.041	0.034	0.090	0.069
	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)
Metro	0.116	0.125	0.055	0.058	0.038	0.040	0.079	0.085
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Nonmetro-metro difference	22%	−15%	13%	−16%	8%	−14%	13.9%	−19%
	(3.55)	(2.76)	(4.06)	(3.33)	(4.73)	(4.11)	(6.30)	(5.12)

Note: Poverty indices are the Watts and Foster-Greer-Thorbecke P_α measures. FMR adjusted are poverty measures after adjusting for spatial-price variation with the FMR index. Nonmetro-metro differences are $(P\alpha_{\text{nonmetro}} - P\alpha_{\text{metro}})/P\alpha_{\text{metro}}$, both using actual levels and FMR-adjusted levels. Standard errors for the poverty measures are estimated following equation 5 using the program described in Jolliffe and Semykina (1999). Standard errors for the differences are second-order approximations by the delta method.

FGT = Foster-Greer-Thorbecke, FMR = Fair Market Rent.

plots the nonmetro-metro percentage differences for the three FGT poverty measures.[23] This panel readily indicates that over all years between 1991 and 2002, all three of the FGT poverty measures indicate that nonmetro poverty is greater than metro poverty.[24] Panel B plots these same differences but for poverty measures that have been adjusted with the FMR index. This panel reveals that the reversal of the relative rankings holds over all years considered. The spatial-price adjusted estimate of nonmetro poverty is lower than the adjusted metro estimate for all FGT measures over all years. This panel indicates that most of the FMR-adjusted nonmetro poverty estimates are 10% to 25% less than the FMR-adjusted metro estimates.

The analysis presented in panel B of Figure 2 is based on the 2001 FMR index for all years. Ideally, one would prefer using an FMR index for each year, but access to the county identifiers in the CPS data is not publicly available, and they are necessary to properly merge files. Use of the 2001 FMR index then implicitly assumes that the spatial distribution of prices has not changed significantly over the years. A cursory look at the early FMR data files suggests that the results are not likely to be qualitatively affected by this simplifying assumption. To assess this, I examined the simple mean FMR for nonmetro counties and mean FMR for Metropolitan Statistical Areas between 1991 and 2004. These estimates will not be comparable to area means of the FMR index, because they have not been population weighted, but significant temporal variation in the estimates would suggest that the findings in Table 1 could be sensitive. Over the 14 years examined, the nonmetro mean FMR was between 68% and 77% of the metro mean FMR.

Age and the FMR-Induced Change in Nonmetro Poverty

Previous research on demographic differences in area poverty rates has indicated that the

23. The relative difference in poverty uses the metro poverty level as the base and can be expressed as $(P\alpha_{\text{nonmetro}} - P\alpha_{\text{metro}})/P\alpha_{\text{metro}}$.

24. Panel A also reveals a primary finding of Jolliffe (2003b)—namely, that nonmetro-metro poverty differences diminish as one considers measures that are sensitive to the income distribution of the poor. In other words, P_0 indicates a much greater nonmetro-metro difference in poverty than does P_1 and P_2.

FIGURE 2
Nonmetro-Metro Poverty Differences from 1991 to 2002

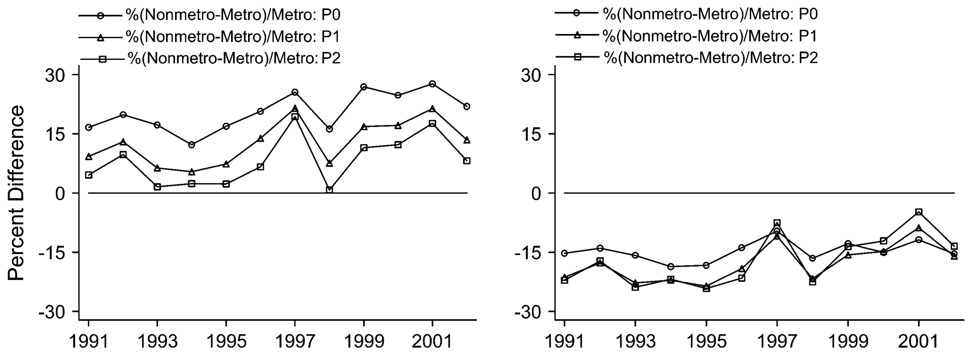

Panel A: P_α Differences (No adjustment) *Panel B: P_α Differences (FMR adjusted)*

Note: In the left panel, the Pα lines plot the difference between nonmetro and metro poverty as measured by P_0, P_1, and P_2 using metro poverty as the base, or $(P\alpha_{nonmetro} - P\alpha_{metro})/P\alpha_{metro}$. In the right panel, the Pα lines are adjusted using the Fair Market Rent index to correct for geographic differences in prices.

nonmetro poor are more likely to be disabled and retired whereas the metro poor are more likely to be going to school. This difference suggests differences in the age structure of the poor in metro and nonmetro areas. Figure 3 provides some support to this conclusion by graphing the age distribution of the poor in metro and nonmetro areas in 2001 and 2002. For both years, the nonmetro age distribution lies above the metro age distribution as age increases (and below for younger ages). There are disproportionately more poor persons over the age of 40 in nonmetro areas than in metro areas (and, similarly, disproportionately more of the metro poor are under 40).

Table 1 indicates that the FMR-adjusted poverty rates produce a stark reversal of the relative rates of nonmetro and metro poverty, Table 2 indicates that age is an important correlate of this readjustment. In both 2001 and

FIGURE 3
Age Distribution of the Poor by Area, 2001 and 2002

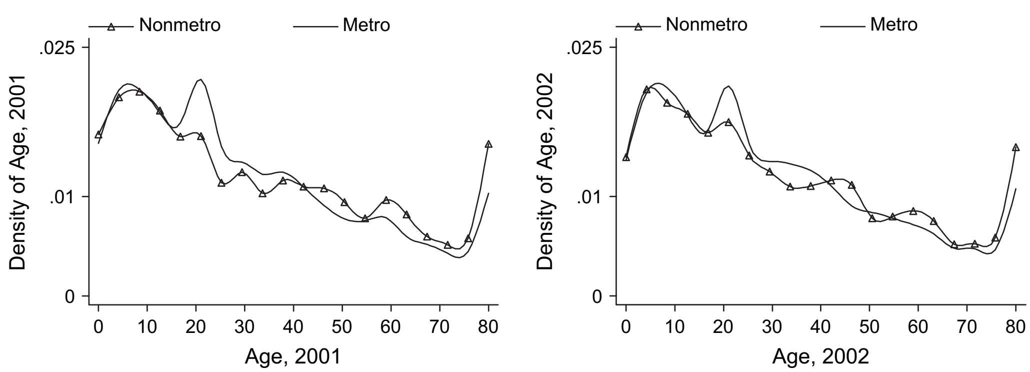

Note: Kernal density estimates of age in years for metro and nonmetro poor persons. Panel on lefthand side is the age density of the poor in 2001, and the righthand side is age density in 2002. The Epanechnikov kernal is used with a smoothing parameter set to 1.

TABLE 2
Average Age of the Poor by Area

	2001		2002	
	Metro	**Nonmetro**	**Metro**	**Nonmetro**
Poor	29.75 (0.257)	32.27 (0.506)	30.20 (0.252)	32.00 (0.484)
FMR poor	30.08 (0.252)	30.59 (0.561)	30.40 (0.245)	30.25 (0.516)

Note: Age in years of the poor by metro and nonmetro residence. FMR poor are those designated as poor after adjusting income with the FMR index. Standard errors corrected for complex design.
FMR = Fair Market Rent.

2002, the nonmetro poor were about two years older than the metro poor on average. Adjusting for cost-of-living differences with the FMR reduces the average age of the nonmetro poor by almost two years, eliminating the difference in average ages.

Figure 4 explores this issue in more detail by plotting the age distribution of the nonmetro poor in 2001. In this figure, the nonmetro age distribution is plotted for those who are poor following the federal definition and using the FMR-adjusted measure. The age distribution for the FMR poor lies below the age distribution of the poor at ages greater than 60 years and above for ages less than about 25 years. This figure reveals that the FMR adjustment reclassifies disproportionately more of the nonmetro elderly from poor to not poor.

FIGURE 4
Age of the Nonmetro Poor, Comparing the Poor with the Fair Market Rent–Adjusted Poor

Notes: Kernal density estimates of age in years for nonmetro poor persons. The unmarked line is the age density for those nonmetro residents who were poor in 2001. The marked line is the age density for those nonmetro residents who are identified as poor using the Fair Market Rent (FMR) index. The Epanechnikov kernal is used for both estimates with a smoothing parameter set to 1.2.

Sensitivity Analysis of the FMR Index

The FMR index currently under consideration by the Census Bureau implicitly assumes that there is no spatial variation in the prices of all goods other than housing costs. To examine the sensitivity of the findings in this article to this assumption, I examine the case where prices of housing and other goods are correlated. Although, to the best of my knowledge, there is no published data on the magnitude of the correlation, I assume that a common belief is that the sign of the correlation coefficient is positive—namely, high-rent areas face high prices in other goods and services. Following this line then, I consider a case where the coefficient of correlation between housing and other prices is 0.2. This implies that areas with housing prices 10% higher than the baseline are also areas with prices of other goods that are 2% higher. With the fixed budget weights, this implies an FMR index that is 5.5% higher than the baseline $(10 \times 0.44) + (2 \times 0.56) = 5.5$. An alternate way of stating this is that the assumption of positive correlation will amplify the spatial variation by approximately 25%.[25]

25. In the example of housing prices being 10% greater than the baseline, the FMR index (with no correlation in prices) is 4.4. When assuming 0.2 positive correlation in prices of housing and other goods, the FMR index increases by 25% to 5.5.

FIGURE 5
Nonmetro-Metro Poverty Differences, Sensitivity Analysis

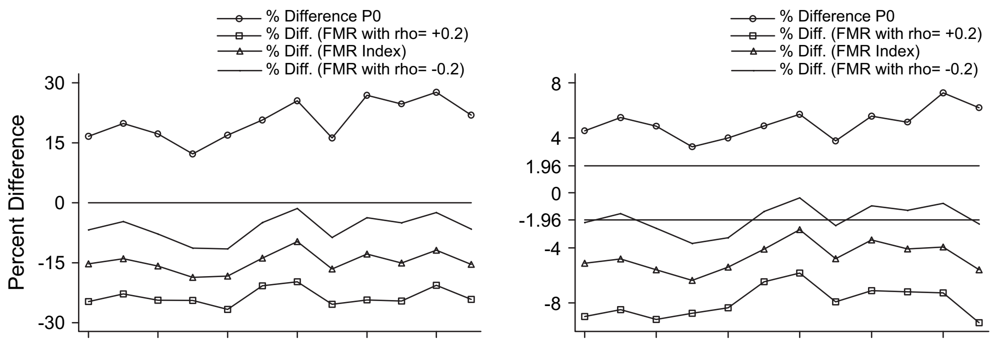

Panel A: P_0 Differences *Panel B: T-stats of P_0 Differences*

Notes: In panel A, the top line plots the difference between nonmetro and metro poverty as measured by P_0 using metro poverty as the base, or $[(P_{0,\text{ nonmetro}} - P_{0,\text{ metro}})/P_{0,\text{ metro}}] \times 100$. The three lines below all reflect the relative difference in P_0 with cost-of-living allowance adjustments. The middle line of the three uses the baseline Fair Market Rent (FMR) index. The line above the baseline assumes that nonrent prices are negatively correlated ($\rho = -0.2$) with FMR, whereas the bottom line assumes positive correlation ($\rho = 0.2$). In the right panel, the lines plot t statistics of the test for whether the nonmetro-metro differences are statistically significant.

For the sake of sensitivity analysis though, I consider the case of negative correlation between prices of housing and all other goods and services. I examine the case where the correlation coefficient is -0.2, which implies a dampening of the variation in prices by 25% (relative to the FMR index). If the assumption of negative correlation is justified, then the FMR index would overstate the true spatial variation in prices.

Panel A of Figure 5 presents the percent difference between the nonmetro and metro prevalence of poverty (P_0) under several assumptions. The single line above zero is this difference without any adjustment for spatial price variation. This line reflects the official federal estimates of poverty in nonmetro and metro areas, indicating that nonmetro poverty has been uniformly higher than the prevalence of poverty in metro areas. A simple average over the 12 years examined indicates that nonmetro poverty has been 20% higher than metro poverty. The three other lines, all falling below zero, have been adjusted using the FMR index assuming negative, positive, and no correlation between prices of housing and other goods. The middle line of the three is the baseline FMR index. The most negative line assumes positive correlation, and the line closest to zero assumes correlation of -0.2.

As expected, positive correlation magnifies the primary findings. The relative ranking of nonmetro and metro poverty is completely reversed. With the coefficient of correlation at -0.2, the findings are dampened, but the estimated level of nonmetro poverty is still lower than metro poverty over all years examined. Panel B reports the t statistics for these relative differences. For the FMR index and the FMR index with 0.2 correlation in prices, all differences are statistically significant. With correlation of –0.2, the difference is statistically significant (at the 5% confidence level) in 6 of the 12 years considered. If the true correlation in prices is negative and greater than 0.2 in magnitude, then the primary finding of a reversal in the rankings would no longer hold.

IV. CONCLUSION

The prevalence of poverty has been greater in nonmetro areas than in metro areas in every year since the 1960s, when poverty rates were first officially recorded; accordingly, federal funds for social safety nets and community development have favored nonmetro areas. The federal government is currently examining experimental poverty measures that, among other changes, adjust poverty rates for spatial cost-of-living differences. Currently, the preferred experimental index is one based on the FMR data, which reflects spatial differences in the rental cost of low-income housing. The purpose of this article is to examine how the use of this index to adjust for cost-of-living differences affects the distribution of poverty across metro and nonmetro areas.

The primary finding is that adjusting poverty rates with the FMR index results in a stark and complete reversal of the nonmetro-metro poverty profile. With no adjustment for cost-of-living differences, the prevalence of poverty is higher in nonmetro than metro areas over the last 12 years. (The depth and severity of poverty are also higher in nonmetro areas, but in about half the cases the differences are not statistically significant.) When the FMR index is used to adjust for cost-of-living differences, the prevalence, depth, and severity of poverty are higher in metro areas than in nonmetro areas over the last 12 years. In 2001, for example, the prevalence of nonmetro poverty was 28% higher than in metro areas. Once adjusted for cost-of-living differences, this is reversed and the prevalence of poverty in nonmetro areas is 12% lower than in metro areas.

Holding the national poverty rate fixed, the analysis also examines how adjusting for cost-of-living differences affects the age composition of the poor. The nonmetro poor consist disproportionately of the elderly population, many of whom are living on fixed incomes near the poverty line. Using the FMR index to adjust for cost-of-living differences results in reidentifying many of these elderly poor as nonpoor. The average age of the nonmetro poor drops from 32.3 years to 30.6 years when adjusting for cost-of-living differences. To the extent that these elderly people are receiving federal funds that are tied to poverty rates, they have the most to lose from this reform. More generally, using the FMR index to adjust poverty rates for cost-of-living differences could potentially have significant adverse affects on funding for nonmetro social safety nets and developmental block grants.

APPENDIX

TABLE A-1

Comparing 90% Confidence Intervals Resulting from Synthetic-Design and Census-Recommended Correction: 1999 Current Population Survey Headcount Poverty Measures

		Estimated 90% Confidence Intervals						
		Direct and Implied Cstimates from the P-60 Report						
Characteristics	**Ratio or Percent Poor**	**Reported Table A**	**Implied by Levels**	***a*, *b* Percentage**	***a*, *b* Ratio**	**Match *a*, *b* Categories**	**Synthetic Design**	**Random Sample**
Persons	11.8	0.3	0.33	**0.33**	*	yes	0.33	0.16
Persons in families	10.2	0.3	0.34	0.34	*	no	0.36	0.17
White	9.8	0.3	0.34	**0.33**	*	yes	0.31	0.16
Black	23.6	1.2	1.20	**1.20**	*	yes	1.24	0.66
Under 18	16.9	0.7	0.65	**0.65**	0.66	yes/no	0.64	0.37
18–64 years	10.0	0.3	0.39	**0.39**	*	no	0.30	0.20
65 years+	9.7	0.5	0.53	**0.53**	0.53	yes	0.53	0.43
Families, total	9.3	0.3	0.33	0.28	**0.34**	yes	0.32	0.29
Total	9.3	0.3	0.33	0.28	**0.34**	yes	0.32	0.29

Note: Confidence intervals are listed in percentage points, and the asterisk (*) denotes that the number is undefined (square root of a negative number). The first four columns of confidence intervals are derived from the Dalaker and Proctor (2000) P-60 report on poverty. The bold estimate marks whether the Census Bureau considers the estimate a percentage or a ratio. The next column lists whether there is a direct match in characteristics between the poverty estimates and those characteristics assigned *a*, *b* coefficients. The estimates from the synthetic cluster approached are listed next, followed by the confidence intervals from assuming that the data are from a weighted, simple random sample.

Source: Jolliffe (2001).

REFERENCES

Citro, C., and R. Michael, eds. *Measuring Poverty: A New Approach.* Washington, DC: National Academy Press, 1995.

Cushing, B., and B. Zheng. "Re-evaluating Differences in Poverty among Central City, Suburban, and Nonmetropolitan Areas of the US." *Applied Economics*, 32(5), 2000, 653–60.

Dalaker, J., and B. Proctor. "Poverty in the United States: 1999." Current Population Report No. P60-210, U.S. Census Bureau, Washington, DC, 2000.

Datt, G., and M. Ravallion. "Growth and Redistribution Components of Changes in Poverty Measures. A Decomposition with Applications to Brazil and India in the 1980s." *Journal of Development Economics*, 38(2), 1992, 275–95.

Dumond, M., B. Hirsch, and D. MacPherson. "Wage Differentials across Labor Markets and Workers: Does Cost of Living Matter?" *Economic Inquiry*, 37(4), 1999, 577–98.

Fisher, G. "From Hunter to Orshansky: An Overview of (Unofficial) Poverty Lines in the United States from 1904 to 1965." Washington, DC: U.S. Census Bureau, 1997, http://www.census.gov/hhes/poverty/povmeas/papers/hstorsp4.html.

Foster, J., J. Greer, and E. Thorbecke. "A Class of Decomposable Poverty Measures." *Econometrica*, 52(3), 1984, 761–65.

Foster, J., and A. Shorrocks. "Subgroup Consistent Poverty Indices." *Econometrica*, 59(3), 1991, 687–709.

Ghelfi, L. *Rural Welfare.* Washington, DC: U.S. Department of Agriculture, Economic Research Service, 2003, http://www.ers.usda.gov/Briefing/IncomePovertyWelfare/ruralwelfare.

Howes, S., and J. O. Lanjouw. "Does Sample Design Matter for Poverty Comparisons." *Review of Income and Wealth*, 44(1), 1998, 99–109.

Jolliffe, D. "Estimating Sampling Variance from the Current Population Survey: A Synthetic Design Approach to Correcting Standard Errors." *Journal of Economic and Social Measurement*, 28(4), 2003a, 239–61.

———. "On the Relative Wellbeing of the Nonmetropolitan Poor: An Examination of Alternate Definitions of Poverty During the 1990s." *Southern Economic Journal*, 70(2), 2003b, 295–311.

Jolliffe, D., G. Datt, and M. Sharma. "Robust Poverty and Inequality Measurement in Egypt: Correcting for Spatial-Price Variation and Sample Design Effects." *Review of Development Economics*, 8(4), 2004, 557–72.

Jolliffe, D., and A. Semykina. "Robust Standard Errors for the Foster-Greer-Thorbecke Class of Poverty Indices: SEPOV." *Stata Technical Bulletin*, 9(51), 1999, 34–36.

Kakwani, N. "On a Class of Poverty Measures." *Econometrica*, 48(2), 1980, 437–46.

———. "Statistical Inference in the Measurement of Poverty." *Review of Economics and Statistics*, 75(4), 1993, 632–39.

Kish, L. *Survey Sampling.* New York: John Wiley & Sons, 1965.

Kokoski, M. "New Research on Interarea Consumer Price Differences." *Monthly Labor Review*, 114(7), 1991, 31–34.

Koo, J., K. Phillips, and F. Sigalla. "Measuring Regional Cost of Living." *Journal of Business and Economic Statistics*, 18(1), 2000, 127–36.

Moulton, B. "Interarea Indexes of the Cost of Shelter Using Hedonic Quality Adjustment Techniques." *Journal of Econometrics*, 68(1), 1995, 181–204.

Nord, M. "Does It Cost Less to Live in Rural Areas? Evidence from New Data on Food Security and Hunger." *Rural Sociology*, 65(1), 2000, 104–25.

Nord, M., and P. Cook. "Do the Proposed Revisions of the Poverty Measure Matter for Rural America?" Economic Research Service Staff Paper No. 9514, U.S. Department of Agriculture, Washington, DC, 1995.

Orshanksy, M. "Counting the Poor: Another Look at the Poverty Profile." *Social Security Bulletin*, 28(1), 1965, 3–29.

Proctor, B., and J. Dalaker. "Poverty in the United States: 2002." Current Population Report No. P60–222, U.S. Census Bureau, Washington, DC, 2003.

Ravallion, M. "Poverty Lines in Theory and Practice." Living Standards Measurement Study Working Paper No. 113, World Bank, Washington, DC, 1998.

Ravallion, M., and B. Bidani. "How Robust Is a Poverty Profile?" *World Bank Economic Review*, 8(1), 1994, 75–102.

Reeder, R. "How Would Rural Areas Fare under Block Grants?" Economic Research Service, Agriculture Information Bulletin No. 724-03, U.S. Department of Agriculture, Washington, DC, 1996.

Reeder, R., S. Calhoun, and F. Bagi. "Federal Funding in the South: Bringing Home the Bacon, but Where's the Beef?" *Review of Regional Studies*, 31(1), 2001, 1–12.

Ruggles, P. *Drawing the Line: Alternative Poverty Measures and Their Implications for Public Policy.* Washington, DC: Urban Institute Press, 1990.

Sen, A. "Poverty: An Ordinal Approach to Measurement." *Econometrica*, 44(2), 1976, 219–31.

Short, K. "Experimental Poverty Measures: 1999." Current Population Reports No. P60–216, U.S. Census Bureau, Washington, DC, 2001a.

———. "Where We Live: Geographic Differences in Poverty Thresholds." Poverty Measurement Working Paper, U.S. Census Bureau, Washington, DC, 2001b.

U.S. Census Bureau. "Current Population Survey: Annual Demographic File, 2000." Inter-University Consortium for Political and Social Research Document No. 6692, U.S. Census Bureau, Washington, DC, 2000.

U.S. Housing and Urban Development. *Fair Market Rents*, 2005, http://www.huduser.org/datasets/fmr.html.

Zheng, B. "Aggregate Poverty Measures." *Journal of Economic Surveys*, 11(2), 1997, 123–62.

———. "Minimum Distribution-Sensitivity, Poverty Aversion, and Poverty Orderings." *Journal of Economic Theory*, 95(1), 2000, 116–37.

———. "Statistical Inference for Poverty Measures with Relative Poverty Lines." *Journal of Econometrics*, 101(2), 2001, 337–56.