

Report on Adjusting Poverty Thresholds for Geographic Price Differences

Edgar O. Olsen*
Department of Economics
University of Virginia
Charlottesville, VA 22904

Prepared for Research Forum on Cost of Living and the Supplemental Poverty Measure
Brookings Institution, April 28, 2011

Report on Adjusting Poverty Thresholds for Geographic Price Differences

The purpose of this report is to evaluate the proposed Census method for adjusting poverty thresholds for geographic price differences, offer empirically implementable alternatives to the Census approach (including the possibility of making no adjustment at all), and suggest future directions for research on geographic adjustment of poverty thresholds. I am in a particularly good position to do it because Paul Carrillo, Dirk Early, and I have recently produced a panel of price indices for housing services, other produced goods, and all produced goods for all U.S. metropolitan areas and the non-metro part of each state from 1982 through 2010 that will be updated every year and is available together with complete documentation at <http://artsandsciences.virginia.edu/economics/facultystaff/eoo.html#price> . The paper that documents the production of these price indices also compares our housing price index with many other alternatives and reports other comparisons between price indices in the literature. I borrow heavily from this paper because the information in it is highly relevant for the matter at hand.

Any thresholds expressed in dollars that attempt to distinguish the poorest people in the country from others should account for geographic price differences. This is true for any reasonable conception of poverty. At the most abstract level, poverty might be viewed as achieving well-being below some level. Given our current state of knowledge, implementing this concept would require the counterfactual assumption that all individuals experience the same level of well-being from the same consumption bundle. This is counterfactual because it implies that all people have the same preferences. That is, faced with the same budget constraints, they would make the same choices. Researchers who work regularly with data on individual households realize that this is far from the truth. However, if we accept this simplification of reality and base poverty thresholds on the estimated preferences of the average person, the minimum income necessary to attain a certain level of well-being at each location depends on the market prices that prevail at this location. It also depends on the characteristics of the location that determine its attractiveness as a place to live. Another conception of poverty that is less abstract and very policy oriented is the inability to purchase minimum amounts of certain goods that some members of society would like to insure that all members attain. Obviously, different people have different views about the specific minimum quantities, and the political system would

almost surely generate official minima between the extremes. Indeed, this has already happened for some types of goods, for example, the minimum housing standards in HUD's low-income housing programs and the SNAP's Thrifty Food Plan. The amount of money required to purchase the official minima is different in different locations due to geographic price differences.

The purpose and conceptual underpinnings of the Supplemental Poverty Measure (SPM) are not clear to me. Its starting point is a certain arbitrary percentile (the thirty third) of the distribution of the market value of goods consumed. If this approach were applied each year in one location populated entirely by renters who did not receive housing or food assistance (and perhaps met other criteria), a third of all people would be categorized as poor each year no matter how the private economy would evolve in the absence of government action or what government policies were pursued. If this characterization of the approach is correct, the purpose of the SPM would not be to measure the change over time in the fraction of the population who achieve a standard of living below a certain level but rather the difference in these fractions across geographic areas and demographic groups. No matter what its purposes are, the developers of the SPM call for adjusting the threshold for geographic price differences.

Carrillo-Early-Olsen Price Indices

The Interagency Technical Working Group (ITWG) on developing the SPM calls for adjusting the poverty threshold for geographical price differences at the level of metropolitan areas and the non-metro part of each state. Since Carrillo, Early, and Olsen (CEO) provide a fairly refined overall consumer price index at this level of geography for each year in a timely manner, it is certainly a candidate for this purpose. It not only is empirically implementable but also has been implemented. This section describes briefly how this price index was produced. Later sections will compare it with the proposed Census method for adjusting poverty thresholds for geographic price differences and other alternatives.

Our general approach is to first produce cross-sectional price indices for a single year 2000 and then use BLS time-series price indices to create the panel. This approach makes it possible to update the panel early each year immediately after the BLS releases its time-series price indices

for the last month of the previous year. Our initial panel covering the period 1982 through 2008 has already been updated through 2010.

Our geographic housing price index for 2000 is based on data on the gross rent and numerous housing, neighborhood, and location characteristics of about 173,000 rental units throughout the United States. The data set contains observations from each metropolitan area and the non-metro part of each state. Its information about the dwelling unit is about as detailed as the American Housing Survey (AHS), the premier data set on this topic. However, our data set also contains information on the census tract of each dwelling unit which makes it possible to append detailed information on the immediate neighborhood of each dwelling unit from the Decennial Census. In this respect, it is much better than the AHS. The American Community Survey contains much less information about each dwelling unit and its neighborhood.

To construct the rental housing price index, we regress the logarithm of gross rent on 122 regressors representing about 70 underlying variables that describe the attributes of the unit, its neighborhood, and contract conditions and 379 dummy variables for different geographic areas. Our table 1 lists the regressors. The estimated coefficients of the geographic dummy variables are used to construct the housing price index. The fit of the hedonic equation was excellent ($R^2 > .8$), and the coefficients used to create the price indices were estimated with considerable precision. The estimated price indexes were consistent with popular views about differences in housing prices. Among the most expensive places to rent an apartment were San Francisco, Boston, New York City and their suburbs. The least expensive places to rent tended to be nonmetropolitan parts of states and small metropolitan areas in the South. The most expensive place to rent (San Francisco) was somewhat more than three times as expensive as the least expensive (nonmetropolitan Missouri).

For most areas, our price index for all goods other than housing is calculated from the price indices for categories of non-housing goods produced each quarter by the Council for Community and Economic Research, formerly the American Chambers of Commerce Research Association (ACCRA). In order to produce a non-housing price index for areas of the United States not covered by their index, we estimate a theoretically-based regression model explaining differences in the composite price index for non-housing goods for areas where it is available and use it to predict a price of other goods for the uncovered areas.

The theoretical model assumes production functions for housing services and other goods that depend on local labor, local land, imported inputs, and inputs whose prices are the same at all locations and allows for the possibility that their constant terms depend on weather. The model also allows for the possibilities that output prices exceed minimum average cost of production due to local regulations and that the unobserved price index for imported inputs depends on the distance to the nearest large metro area. Solving the housing price equation for the unobserved price of land and substituting into the equation for the price of other goods yields a regression model explaining the price index for other goods in terms of the price of housing services, a wage rate for reasonably homogeneous workers, distance to the nearest large metro area, cooling days, heating days, and precipitation. The results of the estimation of this model suggest no misspecification of functional form, heteroskedasticity, or outliers. No deviations between predicted and observed values of the price index for other goods exceed 10 percent and relatively few exceed 5 percent.

Our overall consumer price index for all areas is a weighted average of the price of housing services and other goods, where the weights are the national average expenditure shares of the two composite goods. Empirically, our index is virtually identical to an ideal consumer price index consistent with a simple utility function based on these expenditure shares. Users who prefer other weights can create their own CPI using their preferred weights and our price indices for housing services and other goods.

Finally, we use BLS time-series price indices to create a panel of price indices from our cross-sectional prices. For quite some time, the BLS has produced time-series price indices for groups of goods and all goods combined for specific metropolitan areas and groups of urban areas based on region and population. Almost all of our metropolitan areas fit unambiguously into one of these categories. Seventy nine of our MSA or PMSA are within the 27 BLS metropolitan areas. For our remaining metropolitan areas, we use the BLS price indices for the relevant population size category in its region. Finally, we use the BLS price indices for the smallest population size category in a region for the nonmetropolitan part of each state in that region, except for Alaska and Hawaii. For non-metropolitan Alaska, we use the BLS indices for Anchorage. For non-metropolitan Hawaii, we use their indices for Honolulu. Our paper (Table 7) describes how we handled a few special cases.

The BLS does not disseminate time-series price indices for our categories of goods, namely, shelter and utilities combined and all other goods as a group. With a trivial exception, we apply their methods and weights to produce these indices [U.S. Bureau of Labor Statistics, 2010, Chapter 17].¹ With this exception, our time-series price indices are exactly the same as theirs would be if they had produced indices for these composites. First, we use BLS methods and time-series price indices for shelter and utilities to create a time-series price index for housing in each area. The BLS reports a composite housing price index that includes household furnishing and operations as well as shelter, fuel and utilities. Our housing index does not include household furnishing and operations. Second, we use our housing price index and the BLS price index for all goods to create a time series price index for goods other than housing. Third, we use these two time-series price indices and the overall CPI to inflate and deflate our three cross-sectional price indices.

Initial Proposed Census Method

Renwick (2011) describes the initial proposed Census Bureau method for adjusting poverty thresholds for geographic price differences. One aspect of the method is implementation of the ITWG recommendation to create separate thresholds for renters, homeowners with a mortgage, and homeowners without a mortgage. Because I do not think that this is the best way to account for the differences in the consumption of members of these three groups and it is tangential to accounting for geographic price differences, I focus my comments on the method based on rental outlays. The same issues are involved in all three cases.

The proposed consumer price index used to adjust thresholds is an expenditure-weighted average of price indices for rental housing and all other goods, where housing includes utilities.

Renwick's rental housing price index is the median rent of two-bedroom units with complete kitchen and bathroom facilities based on data from the American Community Survey (ACS). Since almost dwelling units in the U.S. have complete kitchen and bathroom facilities, this effectively controls for only the number of bedrooms. Renwick's method assumes that prices of other goods are the same everywhere.

¹ The BLS does not collect prices every month in all areas. To obtain an annual price index for these areas, they interpolate to obtain price indices for those months where prices are not collected before averaging over the year. We take a simple average of the reported price indices.

As Renwick recognizes, the initial proposed interarea consumer price index leaves considerable room for improvement. Obviously, there are enormous differences in the characteristics and overall desirability of two-bedroom units with complete kitchen and bathroom facilities. With respect to constructing a geographical housing price index, the issue is the extent to which there are differences in the average overall desirability of two-bedroom units across areas. To the extent that there are differences, Renwick's price index will overstate housing prices in places where the overall desirability of two-bedroom units is greatest and understate it in places where it is worst. Furthermore, contrary to Renwick's assumption, prices of other goods are not the same everywhere.

Although the ACS collects only a small set of rudimentary housing characteristics, an obvious improvement in Renwick's method is to use its PUMS to create a housing price index by estimating a hedonic equation explaining the logarithm of gross rent in terms of all of its housing characteristics and dummy variables for different geographical areas. In addition to the number of bedrooms and whether the unit has complete kitchen and bathroom facilities, the variables reported are the number of dwelling units in the structure (8 relevant categories), number of rooms (9 categories), when the respondent moved into the unit (7 categories), and when the structure was built (9 categories).

CEO's results shed light on the extent to which these variables capture differences in the desirability of housing and the differences in the geographic price indices that result from a hedonic equation based on the ACS and their hedonic equation that is based on much more information about the dwelling unit and its neighborhood. To compare their housing price index with feasible alternatives, CEO use the 5-percent PUMS from the 2000 Decennial Census to estimate a hedonic equation explaining the logarithm of gross monthly rent. This data set contains the same housing characteristics as the ACS. (The ACS replaced the decennial census long form.) The R^2 for the hedonic regression underlying CEO's housing price index is about .8. The R^2 for the hedonic based on the data from the decennial census long form is .33. Obviously, CEO's data enables them to account for differences in the desirability of housing to a much greater extent. CEO also compared the housing price indices across areas based on the two estimated hedonic equations (p. 24 and Table 4). The differences are less dramatic. On average, the price index based on the decennial census data is close to proportional to CEO's more refined

index. After scaling the indices to have the same mean across locations, the mean of the absolute percentage deviation of the census-based index from CEO's index is about 8 percent.

Although prices of other goods vary less across areas than housing prices, they are not the same everywhere. CEO's index of the price of other goods is about 39 percent higher in the most than the least expensive location, and their basic housing price index is positively correlated with their composite index of the price of other goods. The correlation coefficient is .76, and a ten percent increase in the price of housing is associated with about a two percent increase in the price of other goods. Based on these results, the assumption that the prices of other goods are the same everywhere leads to a bias in calculating the overall CPI. It results on average in underestimating the CPI in places with the highest housing prices and overestimating it where housing prices are lowest.

Alternative Data Sets

Bettina Aten and her co-authors (2005, 2006, 2008, 2010) have explored alternative methods for using alternative data sets to produce geographic price indices at various levels of geography. Her initial housing price indices were based on the data from the CPI housing sample that underlies the BLS's time-series price indices. Recently, she has used the ACS data for this purpose. Her price indices for other goods are based on the data underlying the time-series CPI. In my view, these are the best alternatives to the data sets that underlie CEO's price indices for constructing interarea price indices, and research should continue on using them for this purpose. Since I have already discussed the shortcomings of the ACS compared with the CEO housing data for constructing geographical housing price indices, this section discusses the strengths and weaknesses of the CPI housing data relative to the ACS for this purpose and the relative merits of the CPI and ACCRA data for constructing price indices for other goods.

One significant weakness of the CPI data for producing any cross-sectional price index for all metropolitan areas and the non-metro part of each state is that it collects data in only 85 urban areas. As a result, price indices for other areas must be predicted. This is also true for CEO's price index for other goods based on the ACCRA data, but to a much lesser extent. ACCRA provided data for 225 of their 380 areas in 2000. This should not discourage further exploration with the CPI data, but it does require careful thought about the best approach to making the

necessary predictions and some method for assessing their likely accuracy.

The CPI housing data is better than the ACS data in some respects. It has somewhat more information than the ACS about the dwelling unit (for example, the presence of central air conditioning), and it is possible to append to its observations information about the neighborhood at a low level of geography from the last decennial census [Moulton, 1995, Table 1]. On the other hand, the ACS has data from all areas and a much larger sample. Furthermore, it has been shown to produce a reasonable approximation of a more refined housing price index. Therefore, it is not clear which forms the better basis for constructing a cross-sectional housing price index for all metropolitan areas and the non-metro part of each state. Perhaps it makes sense to continue to explore each.

Neither the CPI nor ACS housing data is nearly as good as the data used by CEO to produce their housing price index for 2000. For this year, the CEO housing price index is unambiguously better. Furthermore, the CEO price index clearly dominates the interarea housing price index that could be produced with CPI data for other years because the CEO housing price index for other years is based on the superior CEO housing price index for 2000 and the CPI time-series housing price indices at the lowest level of geography that the BLS considers reliable enough to report.

Using the ACS housing data to produce a housing price for years other than 2000 has one advantage over the CEO approach, namely, the ACS contains data for every area in every year. This is not an advantage for the 27 large metro areas for which the BLS produces time-series price indices. The CEO housing price index should be superior to the ACS for these areas that account for a large fraction of the country's population. However, the ACS might have an advantage for other areas. For these other areas, CEO use the CPI time-series housing price index for all metro areas in a region in a given size class (for example, metro areas in the northeast with population between 50,000 and 1,500,000) to inflate or deflate the 2000 housing price index for all urban areas in this category. Since some areas experience larger and others smaller than average increases in housing prices, this leads to overstatements of prices in some areas and understatements in others. The reported price indices for the large metro areas in each region shed some light on the likely divergence in rates of increase in housing prices across metro areas in a category. Across the four large metro areas in the Northeast, the percentage

increase in the annual housing price index from 2000 to 2010 ranged from 24.0 in Boston to 28.6 in New York. Across the eight in the Midwest, it ranged from 17.6 in Detroit to 22.1 in Milwaukee; across the six in the South, it ranged from 15.5 in Atlanta to 30.0 in Tampa; and across the nine in the West, it ranged from 16.8 percent in Phoenix to 31.7 in Honolulu. Whether the divergence in the percentage increase in housing prices across areas in the broad categories offsets the advantages of the CEO data and methodology in other respects is an open question.

Unlike the ACCRA data set, the CPI data is not available to independent researchers. Therefore, CEO could not use it to produce their price index for non-housing goods. For the locations where CPI data is collected, there is little doubt that it would generate more accurate price indices for these goods. The CPI data set is collected by professionals. It also has more individual price observations each year than ACCRA (about 1,000,000 versus 270,000 in recent years) and prices many more goods (about 370 versus 59). Like ACCRA, the CPI data set covers only urban areas. However, the CPI collects data from many fewer urban areas than ACCRA (85 versus more than 300). Therefore, using the CPI data to produce price indices for all metropolitan areas and the non-metro part of each state requires predictions for many more areas with the CPI data, with inevitable prediction errors.

Although the prices of other goods vary much less across areas than housing prices, the accuracy of CEO's price index for other goods based on the ACCRA indices is relevant for assessing their overall consumer price index. Koo, Phillips, and Sigalla (2000) shed light on the reliability of the ACCRA index compared with an overall price index based on CPI data, albeit in a comparison limited to 23 metropolitan areas. Specifically, they compare ACCRA's cost-of-living index with a cost-of-living index based on the CPI data [Kokoski, Cardiff, and Moulton, 1994]. When the same simple formula and expenditure weights are used to produce the cost-of-living indices and the two indices are rescaled to have the same mean, the mean of the absolute percentage deviations between the cost-of-living indices is 5.8 percent. More research on this matter is clearly desirable. However, these results suggest that ACCRA price indices do a reasonably good job capturing cost-of-living differences.

Conclusions

The Carrillo-Early-Olsen interarea price indices will provide the best method for adjusting poverty thresholds for geographic differences in the overall cost of living in the usual manner in the near future. These price indices for a year can be produced as soon as the BLS publishes its time-series price indices for the last month of that year. The CEO indices are unambiguously better than the initial proposed Census method for this purpose. Additional research to produce the best possible geographic consumer price indices using the CPI, ACS, and other data sets should be accelerated. Such price indices have many uses in addition to adjusting poverty thresholds. CEO's purpose was to produce price indices that are better than current practice in dealing with geographic price differences until such time as others with more expertise in the production of price indices, better access to existing data, and more resources for new data collection produce better price indices.

However, two broad alternatives to the standard approach to adjusting for geographical differences should be explored further. One is to price the difference in the cost of specified amounts of a subset of goods that represent societal views about minimally acceptable consumption levels, e.g., SNAP's Thrifty Food Plan and the minimum housing standards in HUD's low-income housing programs. This requires interarea price indices for the goods involved in the basket. A second approach attempts to approximate the amount of money needed for a person with average preferences living in different locations with different prices and amenities to attain the same level of well-being. Based on a simple model that captures many important aspects of reality, nominal labor earnings of full-time workers with the same skills, energy, job, and preferences will vary across geographic areas with different prices of consumer goods and area-wide amenities so that they attain the same level of well-being in all areas in which they live. Therefore, if we want to set poverty thresholds to insure the same level of well-being in all locations, a regression of nominal labor earnings on variables that capture differences in skills and jobs and dummy variables for different geographic areas will produce the relevant adjustment factors. Unlike the proposed Census Bureau adjustment factors, the coefficients of the geographic dummy variables in this regression account for both price and amenity differences. Price indices are not needed to implement this approach. David Albouy and others have already made considerable headway on this approach. Although it is quite different from

conventional thinking about accounting for geographical differences in adjusting poverty thresholds across areas, I think that we should not rush to judgment on this approach. It is very intuitive that the income needed in a locality to attain any level of well-being depends on both the prices and amenities in that area.

References

- Aten, Bettina H. 2005. "Report on Interarea Price Levels." U.S. Department of Commerce, Bureau of Economic Affairs Working Paper 2005-11.
- Aten, Bettina H. 2006. "Interarea Price Levels: An Experimental Methodology." *Monthly Labor Review* 129(9): 47-61.
- Aten, Bettina H. 2008. "Estimates of State and Metropolitan Price Parities for Consumption Goods and Services in the United States, 2005." http://www.bea.gov/papers/pdf/aten_estimates_state_metro_2005.pdf
- Aten, Bettina and Marshall Reinsdorf. 2010. "Comparing the Consistency of Price Parities for Regions of the U.S. in an Economic Approach Framework."
- Carrillo, Paul E., Dirk W. Early, and Edgar O. Olsen. 2010. "A Panel of Price Indices for Housing, Other Goods, and All Goods for All Areas in the United States 1982-2008" <http://artsandsciences.virginia.edu/economics/facultystaff/eoo.html#price>
- Kokoski, Mary F., Patrick Cardiff, and Brent R. Moulton. 1994. "Interarea Price Indices for Consumer Goods and Services: A Hedonic Approach Using CPI Data." U.S. Department of Labor, Bureau of Labor Statistics, Working Paper 256.
- Koo, Jahyeong, Keith R. Phillips, and Fiona D. Sigalla. 2000. "Measuring Regional Cost of Living." *Journal of Business & Economic Statistics*, 18(1): 127-136.
- Moulton, Brent R. 1995. "Interarea Indexes of the Cost of Shelter Using Hedonic Quality Adjustment Techniques," *Journal of Econometrics* 68(1): 181-204.
- Renwick, Trudi. 2011. "Geographic Adjustments of Supplemental Poverty Thresholds: Using the American Community Survey Five-Year Data on Housing Costs."
- U.S. Bureau of Labor Statistics. 2010. *BLS Handbook of Methods*. <http://www.bls.gov/opub/hom/pdf/homch17.pdf>.