Report for Cost of Living and the Supplemental Poverty Measure

By David Albouy, University of Michigan and NBER

Introduction and Summary

The idea behind indexing poverty measures to local costs-of-living has considerable merit. Prices of goods and services vary considerably across regions in the United States, and these generally reflect underlying differences in the cost of local factors, namely land and labor. The evidence suggests that these price differences are real, and do not just reflect differences in the quality of goods across regions. As a result, households in expensive areas cannot buy as many consumption goods as households in cheaper areas with the same nominal income. Therefore, if other circumstances in these cities are otherwise equal, households in more expensive areas are worse off than households in less expensive areas, and poverty is mis-measured without a cost-of-living adjustment.

When measuring poverty across areas, it is fundamental to consider why costs differ across areas in the first place. If regional labor markets are competitive and in equilibrium, and if households are sufficiently mobile, higher-cost areas should offer either higher paying jobs, or nicer amenities. If they only offer higher-paying jobs, then wages should rise in step with local costs, meaning that equally-skilled households should receive the same real incomes across areas, while having higher nominal incomes in higher-cost areas, without being better off. In this case, cost-of-living indexation is valid, as households are more likely to surpass a poverty threshold in high-cost areas, without being better off. Moreover, when redistributive programs are tied to geographically uniform thresholds, households are discouraged from living and working in high-cost areas where their labor is needed, as they are less likely to qualify for those benefits. This can lead to a shortage of less-skilled workers in high-cost places.

On the other hand, cost-of-living indexation will lead to inaccurate poverty measures if cost differences are driven by amenity differences, which improve households' quality of life.

In this case, higher costs reflect indirect purchases of non-market amenities, such as sunny weather, lower crime, or cultural appeal, which are being ignored by cost-of-living measures. Households in high-cost areas have simply decided to consume more non-market amenities relative to market goods and are no worse off than households in low-cost areas with the same nominal income. In fact, if firms offer lower wages in nicer areas, households with the same nominal income may actually be better off in high-cost areas. Households in higher-wage areas may be compensated for local disamenities, and may be more likely to surpass a fixed nominal poverty threshold, even though they are no better off. When poverty thresholds are indexed to cost-of-living, this effect is aggravated as less expensive areas offer a lower quality of life, even though they offer higher wages.

Therefore, the main drawback of indexing poverty measures by local costs of living is its failure to account for how differences in local quality of life affect those costs. The degree of its problems is in proportion to how much amenities influence local costs relative to employment opportunities. As will be substantiated below, I make the following recommendations.

1) It is most appropriate to index poverty thresholds at the metropolitan level. Differences in costs within metro areas, across state borders, should be ignored, as lower costs within a metropolitan area are likely to reflect worse amenities, such as higher crime, or greater commuting costs to work.
2) Costs-of-living for poorer households are best approximated using housing rents, rather than housing values. Poorer households are much more likely to rent, and housing values reflect expectations about the future of a city.
3) For working households, an arguably better measure would index poverty thresholds to local wage levels for low-skilled workers, possibly using current measures of "Pay Relatives," already calculated by the Bureau of Labor Statistics. Higher wage levels not only reflect higher costs but also worse amenities, both of which reduce the welfare of households. While wage indexation is superior to cost-of-living indexation, the latter is still superior to no indexation, as amenity differences account for a relatively small fraction of cost-of-living differences across areas in the United States.
4) An additional complication is introduced by the presence of commuting costs to work. Households who commute long distances to city centers may pay a low rent relative to the income they receive, but incur high commuting costs. This issue could

be solved by indexing poverty thresholds to income at the place of work, rather than by the place of residence.

5) For non-working households, whose income depends on non-labor sources, the situation is less clear as it depends on the mobility of those households and their chance of working again. For households that are likely to return to work soon, the insight from the above holds, and wage indexation is superior. For households that do not plan to return to the labor market but with low moving costs, it may best to ignore local wage and cost differences.

The last two points bring into fore the importance of how commuting costs to work and moving costs across residences affect household decisions. In a world where households are completely immobile, they cannot escape price levels that are out of line with local wages and quality of life. In such a world, it could be sensible to index poverty lines to local cost of living. But the United States was largely founded by a mobile immigrant population, and thus it seems appropriate to assume people will move to better their circumstances. Thus, it would be appropriate to further research the moving costs of households and the extent to which the economic equilibrium conditions, used to formulate the recommendations here, hold in reality.

Sources of Cost-of-Living Differences across Areas and Poverty Thresholds

The predictive power of urban economics is largely predicated on the assumption that factors, such as labor, are mobile across space within a country's borders. Households chose where to live based on the market consumption opportunities an area offers, as well as its non-market amenities. To put this concretely, let the welfare of person i in area j be given by the product of their market consumption, quality of life (QOL), which we will assume is the same for everyone in that area.

$$Welfare(i,j) = Consumption(i,j) \times QOL(j)$$

For working households, the consumption opportunities an area offers depends on the income it offers worker relative to local costs of living. This is given by the identity:

$$Consumption(i,j) = \frac{Income(i,j)}{CostofLiving(j)}$$

$CostofLiving(j)$ is presumed to vary only by area, and not by person. Say we want to develop a poverty threshold for everyone with a level of welfare below $\underline{w}$, then putting these two equations together, we have that household i falls below the poverty line if

$$Welfare(i,j) = \frac{Income(i,j)}{CostofLiving(j)} QOL(j) < \underline{w}$$

If households are mobile, they will not choose to live in areas that offers lower consumption and quality-of-life, areas that offer fewer amenities, must compensate households with greater consumption. This is summarized by an equilibrium mobility condition saying that a representative, or average, household has the same welfare regardless of where they live, so that

$$\overline{w} = \frac{Income(j)}{CostofLiving(j)} QOL(j)$$

$$\rightarrow QOL(j) = \overline{w} \frac{CostofLiving(j)}{Income(j)}$$

$$\rightarrow CostofLiving(j) = \overline{w} \times Income(j) \times QOL(j)$$

where $Income(j)$ is the income a reference household earns in area j, and $\overline{w}$ the reference welfare level nationwide. The second line expresses how in equilibrium, quality-of-life should be reflected by high living costs relative to income levels. The third line expresses that higher costs in an area should reflect higher income potential, or nicer amenities.

Putting our equations together we have that a poverty threshold should look at how the income of household i in city j varies with the reference level of income in city j:

$$\frac{Income(i,j)}{Income(j)} < \frac{\underline{w}}{\overline{w}}$$

This would seem to imply that the poverty threshold should simply depend on local income levels, as would be implied by a local relative income standard.[1] However, the reference income level should depend only on how location influences income through wages, and not the

[1] The Department of Housing and Urban Development (HUD) public housing and rental vouchers programs are fairly unique, using local income levels to determine eligibility while using a local index of "Fair Market Rents" to determine benefits. The income limits are calculated by taking percentages, e.g. 80 percent, of median household incomes in a metropolitan area. No adjustments are made for differences in worker characteristics across cities. In Canada, Low Income Cut-Offs (LICOs), used to calculate poverty and determine eligibility for some programs, increase with the population size of a community.

composition of locals. One should consider how a reference household would be paid differently across metro areas to compensate them for local quality-of-life and cost-of-living differences.

Cost-of-living differences may be illustrated graphically if we approximate them through local housing rents. Housing-rent differences are expressed relative to wage differences for low-skilled workers (High-school diploma or below) across metropolitan areas in Figure 1 below (see the last page), using data from the American Community Survey from 2005 to 2009. Both measures are adjusted for quality differences using hedonic regression equations, controlling for observable differences in workers and housing characteristics, and are expressed in logarithmic differences relative to the national average, which for small numbers are close to percentage differences.[2]

In the graph, we see that wages in Miami, FL are close to the national average, but rents are about a third above the average. Thus, households in Miami are either over-paying to be in Miami, or they are benefitting from the many amenities Miami offers, such as abundant sunshine, beaches, and a vibrant night-life. Meanwhile, wages in Decatur, IL, and Kokomo, IN are close to the national average, but their rents are about a third lower. While tastes certainly differ, this suggests these cities offer less desirable amenities than Miami, but they make up for it by having lower housing costs. If this is true, then using a cost-of-living adjusted poverty threshold is inappropriate across these areas: since income levels are the same in these areas, households making $22,000 in Miami are no worse off than households in Decatur, they simply consume more non-market amenities and fewer market goods.

To see the problem with a fixed poverty threshold, as well as one adjusted for cost-of-living indexation, consider the case of Philadelphia, PA and Grand Junction, CO. Both have the same cost-of-living, although the latter, in its scenic mountainous location, has nicer amenities, as workers are welling to be paid 15-percent less to live and work there. A household may earn $24,000 in Philadelphia and $20,000 in Grand Junction and be equally well off, however in Philadelphia they will be above a poverty threshold of $22,000, while in Grand Junction they would be below it, and thus would only be considered poor in one location. An income-adjusted poverty threshold would move in accordance with these income differences, and could be set at,

---

[2] The percent difference is equal to e^b-1, where b is the coefficient reported. Further detail for the methodology is given in Albouy (2008). Note that these measures cannot control for unobserved quality differences, although there are reasons to believe that those differences are fairly small across metropolitan areas.

say, $25,000 in Philadelphia and $21,000 in Grand Junction, so that the household would be consider poor in either area.

When we consider cities that offer wages and rents away from the average, comparisons of quality of life become more complicated. We see from the dotted regression line that wages and rents are strongly correlated, with a one-percent increase in nominal wages predicting a 2.5-percent increase in rents. These higher nominal wage levels generally occur in larger cities, where workers are more productive. The question becomes whether a rent increase associated with a wage increase lowers the market consumption, or real income, available to a worker.

To determine this, it is important to know the share of household incomes from wages, and the marginal tax rate that they face. This obviously differs remarkably across households, but for poorer households, it appears that about 82 percent of income comes from labor, and that they face an effective marginal tax rate of about 25 percent (including payroll taxes and, state income and sales taxes). Consequently, a one-percent wage increase on average translates to a 0.82 x (1 – 0.25) = 0.65 percent increase in wages. Note that this calculation does not take into account various government transfers, such as Medicaid and housing assistance, which for some households lead to much higher effective marginal tax rates.

The other challenge is to determine how to approximate overall costs-of-living using rent differences alone. For this purpose, I assume households spend a constant share of income on housing (as opposed to a constant amount) of 23 percent. To account for costs-of-living differences due to non-housing consumption, this number is increased to 38 percent, to account for the finding in Albouy (2008) that housing accounts for approximately 60 percent of spatial cost-of-living differences. Thus a one-percent increase in rents is associated with 0.38 percent increase in cost of living, and would be offset by a 0.38 / 0.65 = 0.62-percent increase in wages; inversely, a one-percent increase in wages would be offset by 1.60 increase in the rent level.

Accordingly, the slope of the blue dotted line in Figure 1 is 1.60: metro areas with markers above this line have rents that are high relative to the local wage level, which theoretically are associated with greater amenities. Albouy (2008) documents that this excess-rent measure is strongly associated with climactic and geographic amenities, namely, warmer winters, mild summers, annual sunshine, proximity to coast, and degree of hilliness. However, it is not possible to rule out that other factors may affect these rent levels other than amenities.

Nevertheless, the line implies that cities like San Francisco, Honolulu, and Myrtle Beach all offer an above-average quality-of-life, which seems altogether reasonable.

Cost-of-living adjustments works best for cities along this line: the closer cities are to it, the more cost-of-living differences perfectly offset income differences, implying no difference in quality of life. Overall, the positive relationship suggest that quality-of-life differences account for only a minority of cost differences across cities, meaning that cost-of-living adjustment is likely preferable to no cost-of-living adjustment.

Place of Work, Commuting Costs and Sub-Metropolitan Amenities

Another complication is given by the fact that rents tend to be lower further away from metropolitan centers, where wage levels are the highest. This is explained in the canonical mono-centric city model of Alonso, Muth, and Mills, where households further from the city pay lower rental costs, which are offset with higher commuting costs, both in time and money. Thus, areas further from the city center appear to have lower costs through rents, although households may actually be incurring much higher commuting costs, which are not accounted for by a rent-based index. This is seen in Renwick's cost figures for West Virginia within the Washington, DC Metro Area, which are remarkably lower ($760) than the Metropolitan average ($1,246). Presumably these households are either commuting long distances to get higher-paying jobs in the central Washington, or taking lower-paying jobs closer to them.

If city jobs pay just enough to compensate for the commuting costs, a household will be indifferent between taking a job near or far. If a household chooses to take the central-city job, their higher salary may cause them to exceed a poverty threshold, even if net of commuting costs, they would fall below it. If instead, a household takes a lower-paying job in the exurbs, they are less likely to surpass the threshold. For example, a worker may have to incur an additional commuting cost of $4,000 to take a job in the city, where they get paid $24,000, or take a nearby job for $20,000. In the former case they may surpass a threshold of $22,000, while in the latter they be below it, even though their welfare is no different.

To circumvent this problem, the solution would be to index poverty thresholds by the *wage level at the place of work*, not the place of residence. This would prevent undue

discrimination between households with workers that take distant better-paying jobs from those with workers that take near lower-paying jobs, with lower commuting costs.

A second point concerns areas that are of equal distance to jobs within a metropolitan area, but differ in their local rent levels. For instance, rents in suburban Maryland ($1,224) and Virginia ($1,349) are substantially higher than in the central District of Columbia ($994), despite the fact that housing in the District of Columbia is closer to most of the jobs in the area. These lower costs are almost certainly associated with sub-metropolitan differences in local amenities, such as safety and school quality. Households in central D.C. proper effectively pay part of their rent invisibly by dealing with worse amenities, such as higher crime and lower-quality schools, than in the suburbs. This does not make them better off, and as such, households with the same nominal income do not have higher levels of welfare in lower-rent areas. All of this provides an argument against incorporating state-level differences in cost-of-living within metropolitan areas.

According to the equilibrium theory laid out here, metropolitan cost-of-living adjustment is therefore sensible in so far as they proxy for the income levels offered for those working in them. Indexing poverty thresholds by wage levels at the place of work would provide the best measures of poverty comparable across areas, and would eliminate economic distortions created by means-tested programs using thresholds that discourage households from working in productive central cities. A nationally uniform poverty threshold fails in so far as households take higher paying jobs to compensate for i) local costs-of-living, ii) local disamenities, or iii) greater commuting costs. A threshold adjusted for cost-of-living in a household's residence would be an improvement over a uniform threshold insofar as it approximates an ideal index, which adjusts for the wage levels for where the household members work. My understanding of the data in the United States suggests that this approximation is accurate enough that it would help to mitigate the under-measurement of America's urban poor.

Indexation for Non-Working Households

For households without members that participate in the labor force, the basis for cost-of-living indexation appears to be largely eroded if those households are mobile. Presumably, these households have incomes that are independent of where they live. Those households could

choose to live in cheaper areas, and if they choose to live in a larger, more expensive city, their incomes will not rise so that they surpass as a poverty threshold. Presumably these households pay higher local costs to enjoy the amenities of say, Miami relative to Decatur.

This situation is complicated if the households have strong local attachments or substantial moving costs. For instance, they may be closely tied the location of family members who do work in the area, in which case they are indirectly tied to circumstances in the local labor market. Alternatively, a retiring household may have lived and worked in an area for much of their lives up in a certain area, and have little mobility on retiring. In this case, location decisions are lifetime decisions, as households cannot easily escape to areas where costs are driven less by local wage levels. Ultimately the answer depends on the importance of moving costs and how the importance of local attachments may be treated as a source of welfare to local households. These effects are likely to vary quite substantially across households.

Rents over Housing Prices

My understanding is that poverty thresholds are used in conjunction with income levels to determine poverty status. For good or ill, they ignore household assets, and therefore it seems inconsistent to including housing assets in the formula. This seems to rule out using a different measure of households that own their homes free and clear of payments.

In terms of economic fundamentals, housing values should represent the future stream of rents, net of maintenance and depreciation costs. But poverty thresholds should depend on the current costs of housing, as given by current rents. Therefore it seems inaccurate to include a measure based on housing prices rather than just on housing rents. If a household must spend more to own a house, because it is in an area where rents are expected to appreciate in the coming years, this does not appear to be relevant for a poverty threshold for the current year. If households are mobile, they could choose to sell the house later, and pocket the capital gain. Overall, it seems more sensible and simpler to base a cost-of-living measure based on rent levels alone, as these reflect the current costs to household of living somewhere.

References

Albouy, David (2008) "Are Big Cities Bad Places to Live: Estimating Quality of Life across Metropolitan Areas." NBER Working Paper No. 14472.

Figure 1: Locational Wage and Housing Costs across Areas: 2007

Log Housing-Rent Differential

Log Low-Skilled Wage Differential