Measuring Local Poverty Rates

By

Edward L. Glaeser

Harvard University and NBER

April 26, 2011

I. Introduction

A family with 20,000 dollars in annual earnings can obviously afford far more in Buffalo than in Boston, but economists have still hardly agreed on how to adjust local poverty measurements for local prices of living. A simple procedure might be to just use local price of living indices, but there are at least two reasons why such an adjustment could be highly problematic. First, these price adjustments are based on commodities consumed by the average person and these prices might have little in common with the prices paid for commodities consumed by the poorest Americans. Second, higher prices in some areas might themselves by a reflection of other area amenities—including economic opportunity. Correcting for high prices might therefore lead to an underestimate of the welfare levels in high cost areas.

In this brief essay, I discuss three primary issues associated with the problem of local poverty measurement. I will limit my discussion of many broad topics, such as the basic poverty threshold, and restrict myself to areas related to local price adjustment, housing costs and transportation. I begin in Section II with a discussion of conceptual issues, the spatial equilibrium and local price adjustment. The first point of this section is that it is really impossible to say particularly sensible things about local price adjustment for poverty rate measurement without knowing the uses of the measure. For example, a measure intended to capture the level of human deprivation in an area would be somewhat different than a measure intended to capture the degree of shortage of income, and that would be different from a measure intending to capture the benefits of greater transfer payments in a particular place. A general measure—without defined purpose—is almost impossible to appropriately define.

If the measure is meant to capture human suffering, then price adjustment needs to be limited if those prices are providing an offset for other benefits. If the measure is meant to capture

cash shortfalls, then price adjustments should be more complete. If the measure is meant to reflect the benefits of cash transfers to that area, then the complementarity between other amenities and cash should be taken into account.

The second major theme of that section is that prices faced by the poor and rich can often be quite different and it is quite problematic to focus too much on general prices. In particular, real estate prices for higher quality housing often differ far more across space than the real estate prices for lower quality housing. If we are interested in appropriate poverty rate measures then we need to focus on prices paid by the poor, especially in the area of housing.

Section III then turns specifically to the area of local housing costs and poverty rates. I agree with the recommendation that poverty rates need to be reconsidered for home-owning households— either their income levels need to be raised, to account for the implicit rental income that they are earning, or their housing costs need to be reduced. However, in the case of housing it is far easier to focus on rental units, since housing price appreciation is not involved, and units tend to be somewhat more homogenous.

The most natural approach would be to start with a national housing quantity—not level of expenditure—perhaps a unit with one room per household member. This level could remain fixed over time, which would provide a more time consistent poverty measure, although it could also increase if the goal of the poverty measure was to be more relative. With that fixed quantity, it should be relatively easy to determine the cost of providing that unit in the rental market in different parts of the country. It would also be relatively easy to calculate the price paid for such units by people who were relatively poor, or in areas that were relatively poor, in order to control for housing quality.

But from an urban economist's perspective, housing costs are intrinsically linked to transportation costs. This is the lesson of the Alonso-Muth-Mills model—the canonical model for understanding price differences within a metropolitan area. That link leads me to think that we should be considering including transportation costs within poverty measures in a sensible fashion that doesn't just bump up the poverty rates by increasing the accepted level of expenditures.

Local price adjustment should consider transportation more thoroughly because the ability to get around without a car for every adult is a primary advantage of many higher cost areas, particularly older, denser cities. This advantage gets lost when we don't include transport costs in a poverty measure. Section IV discusses including transport costs in poverty measures. Section V concludes.

II. Conceptual Issues, Spatial Equilibrium and Local Price Adjustments

What is a poverty rate intended to accomplish? The supplemental poverty measures being proposed have no obvious purpose, which makes sense given their experimental nature, but it makes design difficult. I will speculate that there are three eventual roles that such a measure could eventually play. First, the measure might be meant to give people a sense of the overall suffering within the area. Second, the measure might be meant to provide an understanding of the degree of cash shortages. Third, the measure might be meant to capture the marginal benefits of added government transfer payments to low income people in the area.

To move from these three purposes to a discussion of measuring the poverty rate, I will assume that we treat a household as having a well-defined indirect utility function denoted V(Y, P, A), that can be compared across people, where Y represents income, P represents prices and A represents amenities. P and A may both be vectors. Obviously, comparability of welfare across people runs against the core non-comparability disclaimers of basic micro-economic theory, but since those disclaimers are routinely ignored when discussing transfers and redistribution, I will happily ignore them here as well.

The key problem is designing a poverty measure is whether this household is to be considered "Poor." If we were interested in using the poverty rate to measure overall welfare, we would rank someone as poor if and only if V(Y, P, A) falls below some threshold $\underline{V}$. If we were interested in measuring only the cash shortage of this household, we would be interested in knowing whether income, corrected in some way for price, is sufficiently low. If we were interested in understanding the efficacy of cash transfers to this individual, then we would be interested in knowing dV/dY—the marginal utility of an extra dollar given to this household.

Of these three objectives, the second is actually the most difficult to understand conceptually, and that may be why the poverty rate debate is so difficult. It isn't really clear why cash shortfalls are themselves important, if they are telling us either about welfare levels or the marginal benefit of cash transfers. Note that if welfare was separable in amenities so that V(Y, P, A)=W(Y, P)+U(A), then objectives two and three essentially collapse to being the same thing. The marginal utility of an extra dollar would essentially be determined only by Y and P, which together would determine the "real income."

If the goal in measuring poverty is to get some measure of human deprivation, then the goal is to assess whether V(Y, P, A) is below $\underline{V}$. Assume that we knew that the appropriate cutoff in location 0 was Y*(0), then what would be the appropriate cutoff in other locations as the price and amenity level changed? If the price and amenity changes were small, we could use the equality V(Y*, P, A)=$\underline{V}$ differentiate to see how the appropriate income cutoff changes and prices and amenities change. This implies that :

(1) $V_Y dY^* + \sum_j V_{P_j} dP_j + \sum_i V_{A_i} dA_i = 0$ or $dY^* = -\sum_j \frac{V_{P_j}}{V_Y} dP_j - \sum_j \frac{V_{A_i}}{V_Y} dA_i$

In the absence of the amenity terms, we would get that $dY^* = -\sum_j \frac{V_{P_j}}{V_Y} dP_j$, and using Roy's Identity, this yields that $dY^* = -\sum_j X_j dP_j$, which is the standard result underlying price indices—the change in income should equal some base level consumption times the change in price.

But there is also the amenity level issue which seriously complicates the situation. Even if one does not accept the full logic of the spatial equilibrium assumption—the welfare levels are equalized over space—there is still an abundance of evidence suggesting that prices are often higher in places with more attractive amenities such as January temperature. In that case, the amenity levels will counteract the price correction.

If we were considering only individuals with identical human capital, the spatial equilibrium assumption would imply that V(Y, P, A) is equal over space and that $dY = -\sum_j \frac{V_{P_j}}{V_Y} dP_j - \sum_j \frac{V_{A_i}}{V_Y} dA_i$-- notice that this implies that income changes are always offset by changes in prices and amenities. In principle, this logic would imply that welfare is everywhere equal, for a given level of human capital, and the poverty rate should just be found by measuring human capital. Of course, back in the real world, this approach would be wildly implausible because of the inability to accurately measure human capital, the empirical imperfections of the model and the role of uncertainty.

Still, the logic of the model suggests that there are several possible approaches to dealing with amenities. The first would be the direct measurement of amenities and the use of wage and price data to back out the typical valuation for these amenities, as in Roback (1982). A second approach adopts a rule of thumb based on national price and income data that some percent of price increases are typically correcting for amenities and as a result the overall local price correction should be reduced by that amount. A third approach is just to assume that amenities are orthogonal to prices, which is what correcting for prices but not amenities implicitly does.

The framework also helps us to answer the question about the marginal utility of income. In this case, we define Y*(0) not based on the absolute welfare level, but on the marginal utility of income in that area, so that $V_Y(Y^*, A, P) = V_Y^*$. We assume concavity of the implicit utility function so that higher levels of Y will have lower marginal utilities and lower levels of Y will have higher marginal utilities of income. This cutoff is therefore based on where an extra dollar of Federal spending will do the most good.

In order to determine how Y* changes with prices and amenities, we differentiate again and we get: $dY^* = -\sum_j \frac{V_{YP_j}}{V_{YY}} dP_j - \sum_j \frac{V_{YA_i}}{V_{YY}} dA_i$. Even before we get to the issue of correcting for amenities, it is clear that standard price indices are not appropriate. The important thing is the extent that local prices change the marginal utility of extra income. Higher prices tend to reduce the marginal utility of extra income because an extra dollar goes less far. They also tend to increase the marginal utility because people are poorer in real terms. If the goal is to say something about where Federal dollars are better spent, then a naïve cost of living adjustment is not particularly helpful.

The amenity issue is also significant because again there is the high likelihood that amenities and prices will move together. But in this case, the critical concern is not whether amenities are offsetting the utility gains, but rather whether amenities are increasing or reducing the marginal utility of extra income. If extra income raises the marginal utility amenities, so that there is positive sorting on the basis of income across amenity levels, then this will mean that there is a better case for transferring more resources to high amenity areas.

An added issue which is worthwhile emphasizing is that price and budget weights for ordinary consumers are rarely appropriate for lower income consumers. This is not just the issue of core expenditure shares being different across people—the higher share of expenditures on food among the poor—it is also about the nature of the good being quite different at different income levels. In the case of food, the price of restaurants used by the wealthy may differ significantly across space, while the price of groceries may be far less volatile.

The most important area of price heterogeneity, however, is surely in the area of housing. Poor and rich live in different geographic sectors and consume wildly different quality levels of housing. I briefly discuss these issues in the next section.

III. Housing Costs across Space

The Census' American Fact Finder lists Manhattan's median housing unit value at $800,000 over the 2005-2009 period, which is more than four times the national average. Median rent was only 1200 dollars, only 50 percent higher than the national average. New York City's issues are somewhat compromised by rent control, but many pricey areas still seem to have a reasonable amount of lower cost rental units. For example, the median rent in San Jose California, in the heart of Silicon Valley, is 1300 dollars, while prices are $685,000. The issue is not just that housing prices are often quite high relative to rental costs, but that the variation in prices is far higher than the variation in rental costs.

The differences between housing prices and rents reflect several different factors. First, owner-occupied housing also offers the possibility of price appreciation which means that high prices need to be discounted for expected future price gains. Second, owner-occupied housing is quite different typically in size, location and quality from rented housing. Since owner-occupied units are far more common among the rich than among the poor, basing local price measures on housing prices is likely to yield big errors.

The simplest fix for this concern would be to use only rental costs for computing poverty measures. This would eliminate the problems of inferring housing price appreciation and in some sense focus in on a particularly sensible housing option for many poorer Americans. It is possible that some areas, such as Detroit, might have owning options that are cheaper than renting, but the unobserved quality issues there are also quite considerable.

One vexing issue is how to deal with the ownership costs for lower income people who actually do own their own homes. One solution is to determine the income that they could get by selling their property and then investing the proceeds into treasury bonds. That income would then be included in their earnings, and then this could be compared with housing costs for reasonable rental units. This would provide a measure of what they could afford were they to cash out of their homes. This procedure would also imply that rising housing prices would actually be seen as a boon, not a burden, to homeowners, at least if rental costs rose by less.

Another benefit of focusing on rental costs is that quality levels are homogeneous and it would be possible to construct reasonable price indices for two or three bedroom rentals in most major areas. Focusing on a fixed unit size would make intertemporal and geographic comparisons easier, since the poverty rate wouldn't then fluctuate by putting smaller households into bigger units over time.

Housing is a key aspect in local price levels and it can be made into a strong element in a local price index. If rentals of fixed unit size are used as the basis for comparison, the poverty measure can be based on a very clear and appropriate housing measure.

There are some unit quality measures that would need to be addressed, but even Census data contains some key characteristics that could be controlled for. The most important unit characteristic, however, is location and that will differ sharply between rich and poor. This provides another reason to focus on rentals, which will typically be more abundant in areas occupied by the poor, but it creates a big comparison problem: To what extent can we compare the locations of different units in different areas?

This problem is the twin of the amenity problem discussed above. Some amenities, like January temperature, are metropolitan area-wide. Other amenities, like school quality and commutes, are very local. Local prices depend on just such amenities, and if we are trying to control for welfare levels it might make sense to do more to control for local amenities.

Of course, then we have entered the tricky territory of evaluating local governments and other service providers.

The canonical model of local prices—the Alonso, Muth, Mills model—highlighted price differences that came from differences in commuting costs. Some aspects of that model feel out of date today—people no longer commute overwhelmingly to a single city center. But it remains true that different areas will offer different commuting bundles.

Those differences across areas also help explain why the rich and poor live in different parts of metropolitan areas and consume different types of housing. The poor often live in the city center to take advantage of public transportation. Outside of the metropolitan area, they also live in particularly far flung areas, presumably because they are more willing to trade off long commute times for low costs. From the perspective of the Alonso-Muth-Mills model, it seems odd to focus strongly on housing and not focus on transportation.

IV. Transportation?

Typically poverty rates are defined with respect to the costs of food, clothing, shelter and utilities, but these categories omit the second largest expenditure category for most poorer Americans—transportation. In the 19th century U.S. and in many developing countries today, it is reasonable to expect that poorer people would typically walk. In some American areas today, public transportation is a reasonable and often relatively affordable option. In other areas, cars are really the only way to get around. Any serious discussion of living costs for poorer people in different parts of a metropolitan area must consider the different car-related costs associated with living in city centers or suburban areas.

For example, housing units in New York City are more expensive than living in rural Kentucky, but New York City's public transit system is quite extensive. Owning a car would be quite foolish for anyone with less income. In rural Kentucky, the basic activities of life are almost unimaginable without a car. Somehow transportation should be brought into poverty measures in the U.S., to address the issue that transportation costs differ across space and are a serious issue for many poorer Americans.

It would however, be a mistake to simply add car costs to the existing set of expenditures, which would just cause the measured poverty rate to skyrocket. A more sensible approach would bring in transportation costs, in some way that keeps the current poverty rate in some benchmark location constant, perhaps by adjusting other needed expenditures downward. Then the differences in those costs over space and time could be included.

I would be wary also of just using current transportation expenditures among people of any income. Much transportation expenditure has little to do with necessity. A more reasonable

approach might estimate a reasonable number of trips per year, and assume that these will occur by bus or subway if that is available, and if not, the cost of a low end car and insurance will be used instead. Availability can, if necessary, be inferred by a sufficiently low share of people in the relative geographic unit (less than one percent perhaps) taking public transportation to work.

Cars and transportation costs represent 14.7 percent of expenditures for people earning between 15,000 and 19,999 dollars in 2009. This share is lower among the very poor, which are less likely to own a car. Transportation should be brought into poverty figures not just because it is a large expenditure item, but also because it varies a lot in cost over space.

Notably, among car users, the primary differences in costs are going to come from gas prices, which are directly observable, and insurance costs which are also easy to observe. Vehicle maintenance and purchase costs appear to far more constant across space. These facts mean that controlling for differences in car-related costs across space should be relatively simple.

V. Conclusion

Poverty rates play a large role in public debate and they certainly can be improved, but for us to make progress, the most important thing is to ask what we are trying to use them to accomplish. Without knowing what poverty rates are meant to do, it is hard to figure out the right way to address differences in prices across space. Indeed, in many cases higher prices are offsetting higher amenity levels, and this fact means that correcting for higher prices may not be entirely appropriate if the goal is to capture welfare levels.

Local price indices are also unlikely to be well suited for capturing the actual prices faced by the poor. This is true in many areas, but it is particularly true for housing. I strongly support the idea of using rental costs rather than home prices to determine housing costs for the lower income individuals, both because housing quality heterogeneity is less and the problem of expected price appreciation is finessed. Housing is sufficiently well measured that it would be possible to get good measures of prices for a fixed size rental unit over space and time.

While we can measure housing rental costs, it is harder to measure the amenities associated with living in different places. Unobserved amenity differences bedevil local price measurement, and in the case of housing, there is a good case to consider one particularly important amenity—transportation cost differences across space. The fact that some areas require cars while others do not is particularly important for poorer Americans and that calls for integrating transport costs more thoroughly into poverty measurement.

References

Alonso, William (1964) *Location and Land Use*. Cambridge: Harvard University Press.

Mills, Edwin S. (1967) "An Aggregative Model of Resource Allocation in a Metropolitan Area," *American Economic Review 57*(2): 197-210.

Muth, Richard (1969) *Cities and Housing.* Chicago: University of Chicago Press.

Roback, Jennifer (1982). "Wages, Rents, and the Quality of Life", *Journal of Political Economy*, Vol. 90, no. 4 (December 1982): 1257-78.