COST OF LIVING AND THE SUPPLEMENTAL POVERTY MEASURE

A Research Forum

Submitted to the Office of the Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services

James P. Ziliak, Ph.D.
Center for Poverty Research
University of Kentucky

May 11, 2011

On April 28, 2011 the University of Kentucky Center for Poverty Research (UKCPR), in conjunction with the Brookings Institution and U.S. Census Bureau, sponsored a research forum entitled *Cost of Living and the Supplemental Poverty Measure* at the Brookings Institution. David Johnson of Census and Ron Haskins of Brookings served as co-organizers in preparing the forum. Among the more than 60 attendees were representatives from ASPE, Agency for Healthcare Research and Quality, the Bureau of Economic Analysis, Bureau of Labor Statistics, Census Bureau, Congressional Research Service, Government Accountability Office, National Academy of Science, Office of Management and Budget, academia, and think tanks. This brief report provides the rationale and summary of the forum.

The impetus for the forum began in the summer of 2009 when ASPE asked UKCPR to prepare a report on the implications of the Measuring American Poverty Act of 2008 (HR. 6941/S. 3636) on the geographic distribution of poverty, especially in the South. The so-called MAP Act, which was co-sponsored by Representative McDermott and Senator Dodd, proposed adoption of a new poverty measure along the lines suggested by the 1995 report by the National Academy of Science (NAS).[1] Then, in March 2010 the Interagency Technical Working Group on Developing a Supplemental Poverty Measure (SPM) directed the Census Bureau to develop an alternative, experimental measure of poverty that was to be released in the fall 2011 alongside the official poverty measure.[2] Like the measure in the MAP Act, the proposed SPM is based in large part on the recommendations of the 1990s NAS panel on poverty measurement.

One of the most vexing challenges in implementing the supplemental measure is whether and how to adjust the poverty thresholds for differences in cost of living across states and localities. Roughly, the current method proposed by Census for the SPM is to adjust the threshold based on geographic differences in housing costs for renters, owners with a mortgage, and owners without a mortgage for units with two bedrooms across over 400 metropolitan, non-metropolitan, and micropolitan areas using data from the American Community Survey.[3] As housing accounts for about 40 percent of the typical family's budget, the geographic adjustment only applies to 40 percent of the threshold. The assumption then is that cost of living is homogenous across regions on the other 60 percent of the budget.

The implications of a poverty measure that varies across locales based on cost of living differences could have profound effects on fiscal federalism should transfer payments ultimately be tied to such a measure. Indeed, the issue recently gained increased urgency as Section 1416 of the Affordable Care Act of 2010 calls for the Secretary of Health and Human Services to

---

[1] Citro, Constance, and Robert Michael. 1995. *Measuring Poverty: A New Approach*. Washington, D.C.: National Academy Press.

[2] SPM Technical Working Group. 2010. Observations from the Interagency Technical Working Group on Developing a Supplemental Poverty Measure
http://www.census.gov/hhes/www/poverty/SPM_TWGObservations.pdf

[3] Renwick, Trudi. 2010. "Geographic Adjustments of Supplemental Poverty Measure Thresholds: Using the American Community Survey Five-Year Data on Housing."
http://www.census.gov/hhes/povmeas/methodology/supplemental/research/Renwick_SGE2011.pdf

conduct a study on the feasibility and desirability geographic adjustment in the parallel set of federal poverty guidelines used in transfer program eligibility.

Interest in the issue of geographic adjustment of poverty thresholds has been longstanding, including a 1976 HEW report and a 1995 GAO report, in addition to the NAS panel.[4] Beyond recognizing that the issue is important, is complex, and that data needs great, the previous panels had difficulty reaching consensus on how to proceed.

Thus, the goal of the forum was to gather leading economists in a roundtable format (1) to critically evaluate the proposed Census method for geographic adjustment, (2) to offer empirically implementable alternatives to the Census approach (including whether to adjust at all), and (3) to suggest future directions for research on geographic adjustment of poverty thresholds. In order to provide guidance to Census on the SPM, as well as the Secretary as research proceeds on the feasibility of geographic adjustment of poverty guidelines, the ultimate aim was that we would achieve some form of consensus among the experts on whether geographic adjustment was preferable to none at all, and if so, what form that adjustment would take.

Appendix 1 contains the forum agenda. The experts commissioned for the forum were (in alphabetical order):

David Albouy
Assistant Professor of Economics, University of Michigan
Dan Black
Professor and Deputy Dean of Public Policy, University of Chicago
Angus Deaton (Keynote Speaker)
Dwight D. Eisenhower Professor of Economics and Public Affairs, Princeton University
Edward Glaeser
Fred and Eleanor Glimp Professor of Economics, Harvard University
Barry Hirsch
W.J. Usery Chair of the American Workplace, Georgia State University
Edgar Olsen
Professor of Economics, University of Virginia
Stuart Rosenthal
Maxwell Advisory Board Professor of Economics, Syracuse University

In addition to the roundtable discussion and keynote address, presentations were made by Trudi Renwick, Chief of the Poverty Statistics Branch in the Housing and Household Economic Statistics division of the U.S. Census Bureau ; Bettina Aten, Research Economist in the Office of the Director at the Bureau of Economic Analysis in the U.S. Department of Commerce ; and Paul

[4] U.S. Department of Health, Education, and Welfare. 1976. *The Measure of Poverty: A Report to Congress as Mandated by the Education Amendments of 1974* (Washington, D.C.: U.S. Government Printing Office, Apr. 1976); General Accountability Office. 1995. "Poverty Measurement: Adjusting for Geographic Cost of Living Differences." http://www.gao.gov/archive/1995/gg95064.pdf

Sullivan, Research Economist in the Division of Price and Index Number Research of the Bureau of Labor Statistics.

In advance preparation of the forum, the six roundtable discussants were provided a background reading list (see Appendix 2), and each was asked to write a 10 page report addressing the three issues identified above pertaining to the forum goal. Appendix 3 contains the report by each of the roundtable discussants, and Appendix 4 contains the powerpoint presentation from the keynote speaker, Angus Deaton.

Summary

Although the reports represent a diverse set of opinions on the efficacy and challenges of adjusting the poverty thresholds for differences in regional cost of living, during the roundtable discussion the seven experts reached *unanimous* consensus on the following:

> **Recommendation 1**: *Some form of adjustment to the SPM thresholds for geographic differences in cost of living is preferable to no adjustment*.
>
> While recognizing that data limitations prevent the construction of an index that approaches a true measure of cost of living, the panelists agreed that constructing a geographic price index was preferable to no adjustment. In a parallel to the adjustment made across countries for differences in prices, this index might be labeled a Poverty Purchasing Power Parity Index (PPPP) to reflect the fact that the threshold makes an adjustment to incomes of the poor across regions of the United States.
>
> **Recommendation 2:** *The current method of adjusting the SPM threshold for housing price differences across regions but not other components of the consumption bundle is reasonable until better data become available.*
>
> Although housing expenses only account for about 40 percent of the typical household consumption bundle, large-scale and high quality data on the other 60 percent of the consumption bundle at the subnational level do not presently exist. Collecting such data should be a high priority.
>
> **Recommendation 3:** *The adjustment for geographic housing price differences should be based on quality-adjusted rental costs.*
>
> The current experimental approach adopted by Census for the SPM makes a distinction between renters, owners with a mortgage, and owners without a mortgage. The panel believed that only rental cost data should be used because home ownership is an asset (a stock) and the focus of the SPM is on income (a flow). Moreover the panel recommended using quality-adjusted rents from a hedonic regression using data from the American Community Survey (ACS) as the basis for rental housing adjustment. An example of how such an approach could be implemented is provided in the report by Stuart Rosenthal.

Recommendation 4: *New sponsored research to inform how and for whom to adjust thresholds for geographic differences in cost of living should be a high priority.*

The panelists agreed, and indeed were exasperated, at the lack of evidence on certain fundamental issues necessary to justify and implement geographic adjustment. Chief among these is the paucity of evidence on the mobility of the poor, including the frequency of such moves and over what distances. Based on standard models in economics the case for geographic adjustment is strongest when there exist substantial barriers to geographic mobility that prevent the poor from moving to locations that improve well being. How this mobility varies across the working poor versus non-working poor (and within the latter, disabled versus retired) is also crucial to better understand *for whom* to adjust for differences. Another high priority area is research on constructing a geographic price index for a constant-quality basket of goods and services that accounts for the entire consumption bundle. As many of the reports indicate, it is unlikely that full adjustment of the threshold to such an index is optimal, but instead partial adjustment is more likely (e.g. as indicated in the report by Hirsch). But to implement partial adjustment, better measures of constant-quality geographic price indices is a needed first step.

Appendix 1: *Cost of Living and the Supplemental Poverty Measure*

A Research Forum sponsored by
University of Kentucky Center for Poverty Research
United States Census Bureau
The Brookings Institution

Thursday April 28, 2011
The Brookings Institution

10:00 a.m.	Registration and Coffee
10:30 a.m.	Welcome and Introductions: James Ziliak, David Johnson, Donald Oellerich, Ron Haskins
10:45 a.m.	Trudi Renwick, Census Bureau *Geographic Adjustment and the Supplemental Poverty Measure*
11:15 a.m.	Bettina Aten, Bureau of Economic Analysis *Regional Price Parities: 2005-2009*
11:45 a.m.	Paul Sullivan, Bureau of Labor Statistics *Inter-Area Price Indices and Regional Cost of Living*
12:15 p.m.	Lunch
1:00 p.m.	Keynote Address Angus Deaton, Princeton University *Price Indices and the Measurement of Well Being*
2:00 p.m.	Break
2:15 p.m.	Roundtable Discussion Moderator: Ron Haskins, Brookings Institution Discussants: David Albouy,University of Michigan Dan Black, University of Chicago Edward Glaeser, Harvard University Barry Hirsch, Georgia State University Edgar Olsen, University of Virginia Stuart Rosenthal, Syracuse University
4:00 p.m.	Adjourn

Appendix 2: Background Reading for the Research Forum Cost of Living and the Supplemental Poverty Measure

Selected Background Reading on Poverty Thresholds

Aten, Bettina, and Marshall Reinsdorf. 2010. "Comparing the Consistency of Price Parities for Regions of the U.S. in an Economic Approach Framework Bureau of Economic Analysis."
http://www.bea.gov/papers/pdf/comparing_the_consistency_of_price_parities_for_regions_of_the_us_in_an_economic_approach_framework.pdf

Citro, Constance, and Robert Michael. 1995. "Adjusting Poverty Thresholds," Chapter 3 (especially pp. 182-198), In *Measuring Poverty*
http://www.census.gov/hhes/povmeas/methodology/nas/files/adjust.pdf

Garner, Thesia. 2010. "Supplemental Poverty Measure Thresholds: Laying the Foundation," BLS.
http://www.census.gov/hhes/povmeas/methodology/supplemental/research/ASSAGarner%20Poverty%20Thresholds%20paper%2012-29-10.pdf

General Accountability Office. 1995. "Poverty Measurement: Adjusting for Geographic Cost of Living Differences."
http://www.gao.gov/archive/1995/gg95064.pdf

Renwick, Trudi Renwick. 2010. "Geographic Adjustments of Supplemental Poverty Measure Thresholds: Using the American Community Survey Five-Year Data on Housing."
http://www.census.gov/hhes/povmeas/methodology/supplemental/research/Renwick_SGE2011.pdf

SPM Technical Working Group. 2010. *Observations from the Interagency Technical Working Group on Developing a Supplemental Poverty Measure*
http://www.census.gov/hhes/www/poverty/SPM_TWGObservations.pdf

Selected Background Reading on Price Indices

Albouy, David. 2010. "What Are Cities Worth? Land Rents, Local Productivity, and the Capitalization of Amenity Values." NBER WP 14981.

http://www-personal.umich.edu/~albouy/Cityvalue/cityvalue.pdf

Black, Dan, Natalia Kolesnikova, and Lowell Taylor. 2009. "Earnings Functions when Wages and Prices vary by Location." *Journal of Labor Economics* 27:1, 21-47. http://gatton.uky.edu/Faculty/Ziliak/Blacketal_JOLE2009.pdf

Carrillo, Paul, Dirk Early, and Edgar Olsen. 2010. "A Panel of Price Indices for Housing, Other Goods, and All Goods for All Areas in the United States 1982-2008." http://www.virginia.edu/economics/Workshops/olsen/CEOFinalDecember2010.pdf (See also links to related resources on Ed's website http://artsandsciences.virginia.edu/economics/facultystaff/eoo.html)

Chen Yong, and Stuart Rosenthal. 2008. "Local Amenities and Life-cycle Migration: Do People Move for Jobs or Fun?" *Journal of Urban Economics* 64: 519-537. http://www.sciencedirect.com/science?_ob=MImg&_imagekey=B6WMG-4SP3SKN-1-1&_cdi=6934&_user=783137&_orig=search&_coverDate=11/30/2008&_sk=999359996&view=c&wchp=dGLzVtz-zSkzS&md5=134573edae09c6abedd9bbe735a69798&ie=/sdarticle.pdf

Deaton, Angus. 2010. "Price Indices, Inequality, and the Measurement of World Poverty." *American Economic Review* 100:1, 5-34. http://www.princeton.edu/~deaton/downloads/deaton_price_indexes_inequality_and_the_measurement_of_world_povetry_aer_2010.pdf

Deaton, Angus, and Olivier Dupriez. Forthcoming. "Purchasing Power Exchange Rates for the Global Poor." *American Economic Journal: Applied* http://www.princeton.edu/~deaton/downloads/Purchasing_power_parity_exchange_rates_for_global_poor_Nov11.pdf

DuMond, J. Michael, Barry Hirsch, and David MacPherson. 1999. Wage Differnetials across Labor Markets and Workers: Does Cost of Living Matter?" *Economic Inquiry* 37:4, 577-598. http://gatton.uky.edu/Faculty/Ziliak/DuMond_Hirsch_MacPherson_EI_published_1999.pdf

Glaeser, Edward. 1998. "Should Transfer Payments be Indexed to Local Price Levels?" *Regional Science and Urban Economics* 28:1, 1-20. http://gatton.uky.edu/Faculty/Ziliak/Glaeser_RSUE1998.pdf

Glaeser, Edward. 2011. "What Democrats Might Learn from the Census." *Economix Blog, New York Times*. http://economix.blogs.nytimes.com/2011/01/11/what-democrats-might-learn-from-the-census/

Moretti, Enrico. 2010. "Real Wage Inequality." http://emlab.berkeley.edu/~moretti/inequality.pdf

Winters, John. 2009. "Wage and Prices: Are Workers Fully Compensated for Cost of Living Differences?" *Regional Science and Urban Economics* 39, 632-643. http://gatton.uky.edu/Faculty/Ziliak/Winters_RegionalScience2009.pdf

Appendix 3: Reports of Experts Prepared for Cost of Living and the Supplemental Poverty Measure

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 $0.82 \times (1 - 0.25) = 0.65$ 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

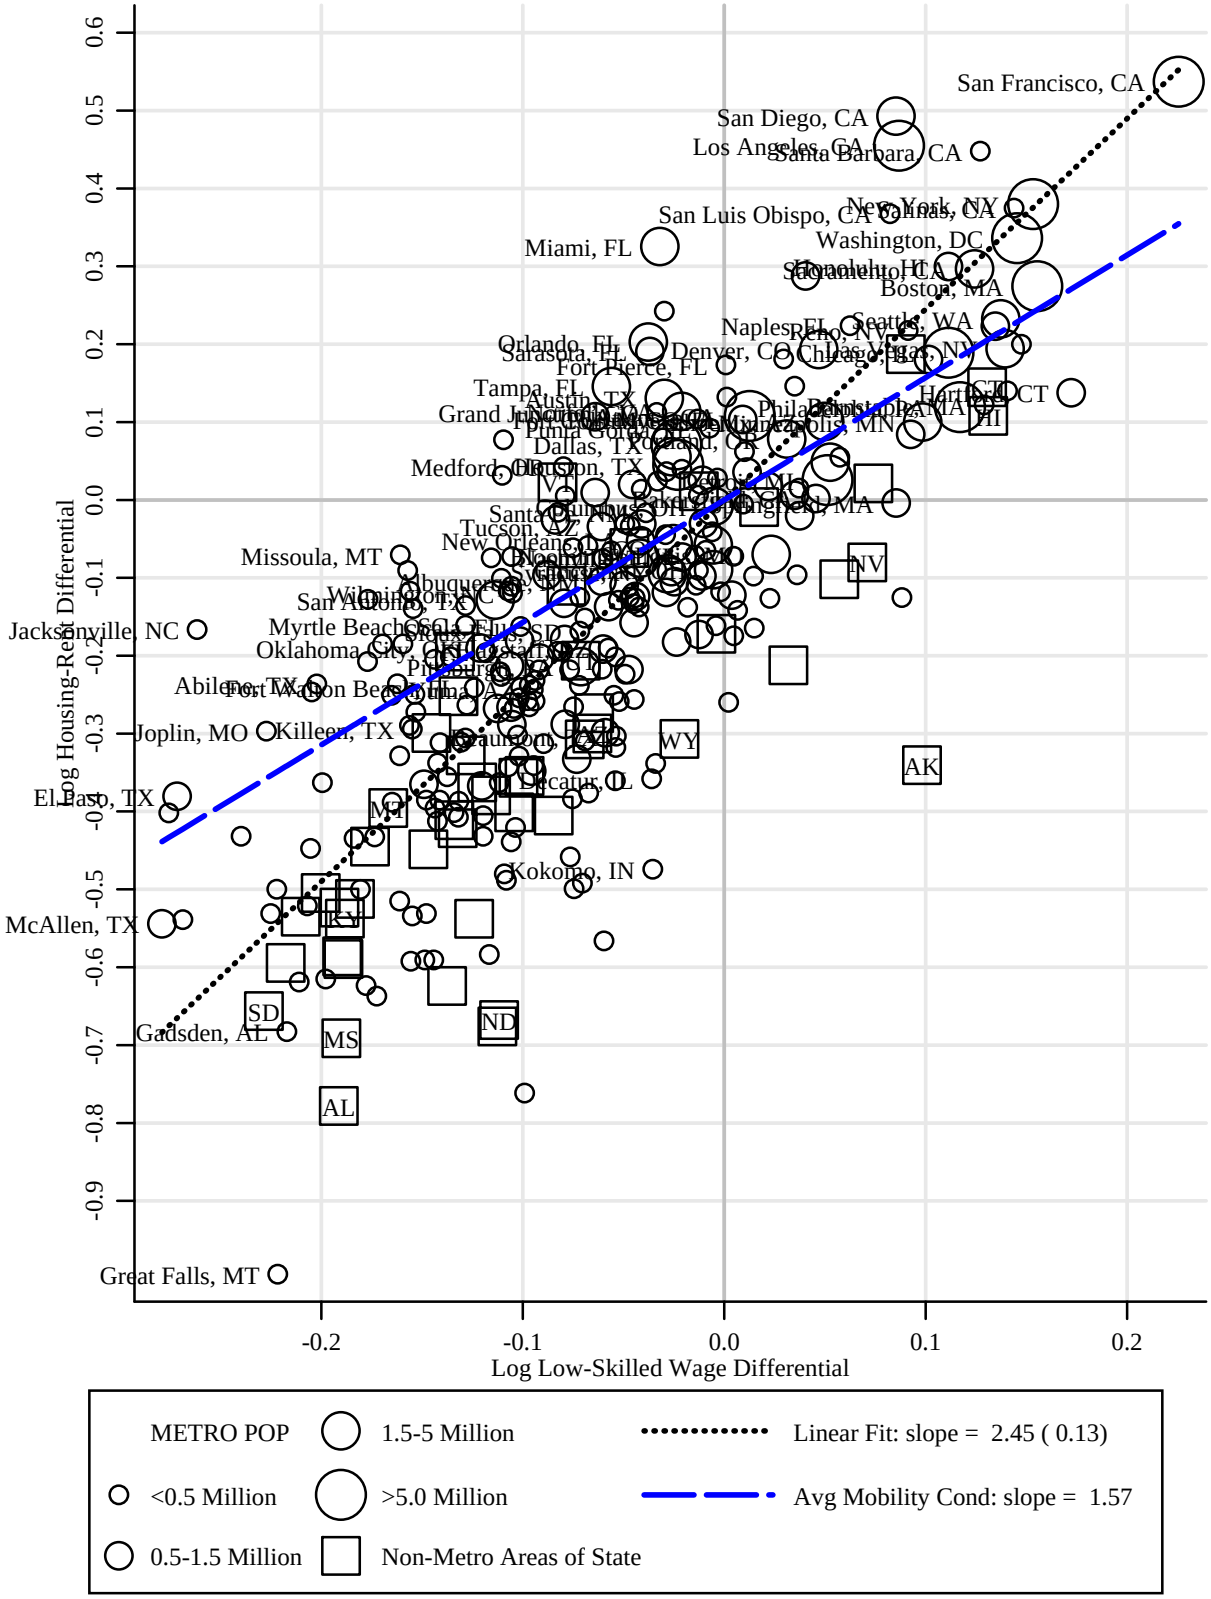

Four Unresolved Issues in the Construction of Regional Cost-of-Living Indices

Dan A. Black
University of Chicago, NORC, and IZA

April 2011

Introduction

The federal government allocates billions of dollars to consumers and firms, often treating the allocations as if agents resided in locations facing the same prices. There are, of course, massive differences in price levels between San Francisco and, say, Buffalo. This strikes some as unfair. The poor in San Francisco will pay much more for food and shelter than the poor in Buffalo. Should we adjust transfer payments to reflect such cost-of-living differences? When federal government measures poverty, the same thresholds are used for San Francisco as for Buffalo. Should we adjust the thresholds of the poverty line to account for differences in the cost of living?

Obviously, there is a reason the Bay area is so expensive; it is a delightful place to live. With its natural beauty and mild climate, residing in San Francisco offers many benefits that are not available to residents of Buffalo. The traditional approach to price indices – allowing agents to purchase the same bundle of goods and services – seems unfair to residents that do not have access to the many amenities of San Francisco.

For economists, there is a natural approach to this question. One would simply ask: "By what factor must I adjust the expenditures in San Francisco to leave the agent equally well-off as an identical agent in Buffalo?" As we shall see, however, the construction of such an index is a formidable problem.

In the next section, I lay out my version of the canonical model of regional differences in prices and wages. In the following section, I discuss four unresolved problems that the model presents for the indexation of the poverty thresholds or transfer payments. In the final section, I discuss whether the index should be used in the indexation transfer payments.

The Basic Haurin-Roback Model of Amenities and Prices

The basic model of location differences in modern urban economics is due to Haurin (1980) and Roback (1982). In this model, locations differ in amenities, wages, and prices. To begin our discussion, divide

all consumption goods into two vectors: those with common prices across locations, which we term traded goods, and those goods that that have prices that can differ across locations. We call these goods "local" goods reflecting the fact that they cannot be traded across sites. Local goods include services such as haircuts and restaurant meals and goods such as housing and utilities that must be consumed locally. Traded goods will include the class of goods that easily shipped to various locations and so prices are approximately the same.

Let x denote the vector of traded goods and let p_x denote their prices. Similarly, let z denote the vector of local goods with prices $p_{z,j}$ where j indexes the location. If we think of consumers as living one period, we have a simple static consumer's problem with a modest complication: Consumers must pick where to live. To keep the notation simple, suppose that there are two locations, or $j \in \{0,1\}$. The two budget constraints facing the consumer are just

$$w_0 = p_x x_0 + p_{z,0} z_0 \tag{1}$$

$$w_1 = p_x x_1 + p_{z,1} z_1, \tag{2}$$

where w_j is the wage paid in the *jth* city and I assume that consumers sell one unit of labor in either city.

Haurin and Roback were concerned that some cities offer better amenities than others, and some cities afford more productive resources for firms. I will ignore these productivity differences (although they may well show up in the equilibrium prices) and focus instead on the amenities that consumers value directly. Toward that end, assume that the consumer has the following utility functions across the two locations

$$U_0 = A_0 f(x_0, z_0) \tag{3}$$

$$U_1 = A_1 f(x_1, z_1), \tag{4}$$

where A_j is the (utility) value of amenities in the jth city and $f(\cdot)$ is the (direct) utility function of the consumer. I have already made a strong restriction on preferences in equations (3) and (4) by making the amenity levels separable from the commodities, but this assumption helps keep this intractable problem a bit simpler.

The consumer, when deciding where to live, solves the utility maximization problem in both cities. Thus, we may express the solutions to these utility maximization problems in terms of the indirect utility functions or

$$U_0 = A_0 h(w_0, p_x, p_{z,0}) \tag{3}$$

$$U_1 = A_1 h(w_1, p_x, p_{z,1}) \tag{4}$$

where $h(\cdot)$ is the indirect utility function. Among the many insights of Haurin and Roback, two are extremely relevant for our problem: (1) Competitive forces will require utility level to be equilibrated between the two cities so that $U_0 = U_1$, and (2) the mechanism that achieves this equilibration will be the capitalization of the amenity values in both the wages and the prices of the local goods.

Assuming that $h(\cdot)$ is strictly increasing in income, we may invert our indirect utility functions to obtain the corresponding expenditure functions or

$$w_0 = e(U_0/A_0, p_x, p_{z,0}) \tag{3}$$

$$w_1 = e(U_1/A_1, p_x, p_{z,1}). \tag{4}$$

Assuming that $U_0 = U_1 = U$, we may now define the regional cost-of-living index, $K_R(U, A_1, A_0, p_x, p_{z,0}, p_{z,1})$, to be just

$$K_R(U, A_1, A_0, p_x, p_{z,0}, p_{z,1}) = \frac{e(U/A_1, p_x, p_{z,1})}{e(U/A_0, p_x, p_{z,0})}. \tag{5}$$

The index $K_R(\cdot)$ is what Samuelson and Swamy (1974) term an "economic index" because it relies on the consumers' preferences.

The regional cost-of-living index answer the question: "By what factor must I adjust income in city one relative to the base city (city $j = 0$) in order to leave the consumer indifferent between the two locations?" As such, the index provides the theoretically correct adjustment for a host of economic problems including the indexation of government transfer payments across disparate sites, the measurement of cost-of-living differences across different cities, the indexation of taxes burdens across location (see Albouy, 2009), or how to define poverty levels in a way that reflects the differences in amenities and prices across locations. [1] For workers who are employed in competitive markets, the Haurin-Roback model suggests that we need only compare the earnings of identical workers in order to know the value of their regional cost-of-living index. Similarly, one could adjust poverty thresholds by, say, the ratio of earnings for college educated workers in the two cities, except for the fact that regional cost-of-living index depends on the utility levels of the relevant consumers.

An important question, therefore, is "When is the regional cost-of-living index independent of the utility level?" It turns out, that if preferences are homothetic, we may express the expenditure function as $^{U}\!/_{A_j}\,\phi(p_x, p_{z,j})$ so that the regional cost-of-living index may be expressed as

$$K_R(A_1, A_0, p_x, p_{z,0}, p_{z,1}) = \frac{A_1^{-1}\phi(p_x, p_{z,1})}{A_0^{-1}\phi(p_x, p_{z,0})}. \tag{6}$$

Pollak (1989) shows that homothetic preferences are both sufficient and necessary for this index to be independent of the value of utility level. The intuition for this result is simple. Because all income elasticities are equal to one when preferences are homothetic, the chosen level of utility does not

[1] This index is directly analogous to the Konüs (1939), or true, cost-of-living index that is theoretical exact measure of the cost of living between two time periods. As a result, the properties of Konus index are also valid properties of the regional cost-of-living index.

matter because the only difference between the very wealthy and the very poor is the scale of consumption.

The assumption of homothetic preferences would, therefore, allow us to construct *the* regional cost-of-living index from information on earnings levels in competitive markets; see Black, Kolesnikova, and Taylor (2009). This would avoid the problem having to measure variations in local goods' prices, in amenities, and in consumption bundles. Unfortunately, there is a consensus within economics – and there are precious few consensuses in economics – that preferences are not, in fact, homothetic. Income elasticities are not uniformly equal to one. This, as we shall see, greatly increases the complexity of the implementation of the regional cost-of-living indices.

Empirical Implementation of the Regional Cost-of-Living Index

When faced with implementing a cost-of-living index to measure the price changes over time, the U.S. Department of Labor's Bureau of Labor Statistics shifts from a cost-of-living index to a "price index." A price index specifies a set of goods and services to be purchased and then proceeds to measure the cost of purchasing that bundle of goods. A popular price index is the Laspeyres type price index, which may be defined as

$$L(p_1, p_0, x_0) = \frac{p_1 x_0}{p_0 x_0}. \tag{7}$$

If the bundle x_0 is chosen to be the consumption in the base period, then we know that the Laspeyres index will overstate changes in the cost of living between period one and the base period zero because it fails to account for the consumers' substitutions among commodities as relative prices change. The best available evidence is that this substitution bias is relatively modest – Braithwait (1980) estimates it to be about 0.4% a year – so ignoring this bias is palatable.

Unfortunately, we have no such luxury when comparing the consumption behavior of agents across cities. Consumers in Buffalo and San Francisco face much different prices for housing, and we

know that consumers with the same utility level will respond to relative prices by selecting different consumption bundles in these two cities. Consumers in Buffalo should consume more housing than consumers in San Francisco because housing is dear in San Francisco. Hence, the first unresolved problem:

Problem 1: When prices differ across locations, how does one select different consumption bundles to reflect differences in relative prices?

The issue, however, is further complicated because the amenity level in San Francisco may be higher than the amenity level in Buffalo. Most people prefer mild climate of the Bay Area to the bitter cold and snow of Buffalo. The Buffalo resident should consume more of the local goods and the traded goods in order to leave the resident indifferent between living in Buffalo and San Francisco. This gives rise to our second unresolved problem:

Problem 2: When amenities differ across locations, how does one select different consumption bundles across locations to compensate for amenity differences?

Our next unresolved issue is one that plagues standard intertemporal cost-of-living indices as well. As I noted above, in the absence of homothetic preferences, the cost-of-living index given in equation (5) is a function of the level of utility chosen to make the comparison. Thus, the cost-of-living index for the very poor may be much quite different than the cost-of-living index for middle class. Again, the issue comes down to the consumption bundles selected. The poor may spend a greater fraction of their income on housing and food than the middle class; hence, higher prices for housing may affect the poor more than the middle class. Our third unresolved, therefore, is:

Problem 3: In the absence of homothetic preferences, how do we select the base utility level?

Of course, my version of the Haurin-Roback model is quite simple. It presumes individuals value all cities equivalently.[2] But a quick look at migration patterns of people, however, shows a remarkable tendency for staying in the same location. For instance, in the Public Use Micro Sample of the 2000 Census, 61 percent of Americans whose highest education attainment is a high school degree report residing in their state of birth. Moreover, for each state in the United States, a majority still reside in the state of their birth. This lack of migration has led many to posit a heterogeneous preference for residing in one's "home town" or at least in close proximity. One might think this arises because people value living close to family and old friends, but it also may reflect that people have city-specific human capital that affects their value of the city. Toward that end, now assume that

$$A_{j,i} = a_j + \varepsilon_{j,i} \tag{8}$$

where a_j reflects the common evaluation of the jth city's amenities and $\varepsilon_{j,i}$ is the ith person's idiosyncratic valuation of the amenities in the jth city. This modest step toward realism greatly complicates the construction of our cost-of-living index.

To see why, consider the value of the cost-of-living index required to leave the ith person indifferent between our two locations. This, of course, is just a slight modification of equation (5):

$$K_R(U, \alpha_1, \alpha_0, \varepsilon_{1,i}, \varepsilon_{0,i}, p_x, p_{z,0}, p_{z,1}) = \frac{e(U/(\alpha_1+\varepsilon_{1,i}), p_x, p_{z,1})}{e(U/(\alpha_0+\varepsilon_{0,i})_0, p_x, p_{z,0})}. \tag{9}$$

While only a slight modification, it presents very different identification and estimation problems. The individuals' valuations of the amenities associated with the location where they live are now idiosyncratic. One might expect, for instance, that the current residents of Detroit in 2011 – given the numerous and severe economic shocks to the city – are people with relatively strong idiosyncratic taste

[2] It also presumes that consumers have identical preferences for goods and services, which is quite unrealistic as well. As this is problem also affecting standard cost-of-living indices, I will ignore the complexities associated with relaxing this assumption.

for living in Detroit. In contrast, in 1961 when the Detroit economy was booming, presumably many of the residents were people with much smaller idiosyncratic taste for living in Detroit.

Problem 4: How do we adjust the cost-of-living index for changes in the idiosyncratic tastes for locations?

Discussion

In the last section, I described several problems that exist in the implementation of regional cost-of-living indices. In this section, I ignore the difficult issues in constructing such an index, and I pose a different question: Should we use the index to adjust transfer payments for differences in the "cost-of-living" index? The largest federal transfer payments are already indexed to local prices, at least implicitly. The OASDI benefits are indexed to the earnings of participants so that residents of San Francisco will receive otherwise higher payments than identical residents of Buffalo because nominal wages are higher in San Francisco than in Buffalo. Because the benefit schedule is nonlinear, the exact nature of the implicit index is difficult to ascertain.

Upon retirement, many seek out locations where the weather is warm and prices are low. Should we seek to reduce the social security payments of retirees who accumulated large benefits by selling their labor services in San Francisco who then exploit the low prices offered on the Gulf Coast of Florida? Or, should we encourage this behavior?

Any payment is implicitly indexed. Offering the same cash value or the same level of medical services in San Francisco as Buffalo are implicitly constructing indices, albeit very different ones. One naïve view of this issue is that individuals in San Francisco should receive higher payments than residents of Buffalo (or residents of rural Appalachia) to compensate them for higher prices they face. This, of course, ignores the fact that the poor on San Francisco have received benefits from their location that have been capitalized into prices. It also presumes, however, that it is economically efficient for the poor to consume the amenities of the Bay area.

Let me offer some evidence to the contrary. In a paper with Gates, Sanders, and Taylor (2002), I argue that the concentration of gay men in San Francisco and other cities with high property prices reflects the fact that the demand for amenities of many cities is luxury good: the rich (or those with relatively modest housing demand) will prefer the expensive cities more than the poor.

Similarly, in Black, Kolesnikova, and Taylor (2009) we note that the returns to schooling in San Francisco are much lower than the returns to schooling in such low-housing price cities as Pittsburgh and Houston. For instance, in 1990, men with a bachelor's degree in San Francisco earned 0.38 log points more than men with a high school degree. In contrast, this differential was 0.53 log points in Pittsburgh and 0.54 log points in Houston. We demonstrate that such differentials are consistent with the simple Haurin-Roback model with housing being a necessity and the demand for locational amenities being a luxury. While this issue is far from settled, it seems wrongheaded to attempt to undo the incentives that markets have created without careful study.

That the labor market solves this formidable problem in such a decentralized environment is remarkable; it is one of the many great accomplishments of Adam Smith's Invisible Hand. It is instructive to think about how the market does it. Firms, of course, do not like paying higher wages so when local economic conditions deteriorate firms limit their wage compensation. The housing market and other locally traded goods will generally see a fall in relative prices, however, that mitigates the decline in consumption that the fall in wages would generate.

Similarly, in booming markets, firms increase wages to retain their workforce and recruit new workers. Again, the housing market will reduce some of this increased compensation through increased prices. These price increases, both for labor and for local goods, are probably contributing to economic efficiency. To the extent that housing a worker in San Francisco is more expensive than Buffalo, we want to provide firms with the correct incentives to consider Buffalo as a location for their production

and provide firms choosing to locate in San Francisco the appropriate incentives to economize on their use of labor.

What is far from obvious is that we want government encouraging transfer payment recipients to pick the Bay Area over Buffalo. Compensating residents of San Francisco for the higher prices of the Bay Area is undoing the incentives for relocation that the market has provided. While one would not want to arbitrarily rule out such compensation, I would want to see a coherent, sensible model that would justify such compensation.

The danger of adjusting the poverty thresholds is that many will view this as "evidence" of the need for such indexation of transfer payments. Currently, the Census Bureau plans to adjust the poverty thresholds largely by differences in the price of shelter. In my view, this is fatally flawed. To see why, consider the following thought experiment: You are guaranteed the same consumption bundle in both San Francisco and Buffalo, but you get to choose where you consume. Tastes differ, but I expect that I will see most of you in the City by the Bay.

References

Albouy, David Y, 2009. "The Unequal Geographic Burden of Federal Taxes" *Journal of Political Economy* 117(4) 635-667.

Black, Dan A., Gary Gates, Seth Sanders, and Lowell Taylor, 2002. "Why Do Gay Men Live in San Francisco?" *Journal of Urban Economics* January 51(1) 54-76

Black, Dan A., Natalia Kolesnikova, and Lowell Taylor, 2009. "Earnings Functions when Wages and Prices Vary by Location" *Journal of Labor Economics* 27(1) 21-48.

Braithwait, Steven D., 1980. "Substitution Bias of the Laspeyres Price Index: An Analysis Using Estimated Cost-of-Living Indexes" *American Economic Review* 70(1) 64-77.

Haurin, Donald R., 1980. "The Regional distribution of population, migration, and climate" *Quarterly Journal of Economics* 95(2) 293-308.

Konüs, A. A., 1939. "The Problem of the True Index of the Cost of Living" *Econometrica* 7(1) 10-29.

Pollak, Robert A., 1989. *The Theory of the Cost of Living Index* Oxford: Oxford University Press.

Roback, Jennifer, 1982. "Wages, rents, and the quality of life" *Journal of Political Economy* 90(6) 1257-1278.

Samuelson, Paul A. and S. Swamy, 1974. "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis" *American Economic Review* 64(4) 566-593.

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.

Barry Hirsch
Andrew Young School of Policy Studies
Georgia State University
April 22, 2011
Revision, May 10, 2011

Adjusting Poverty Thresholds When Area Prices Differ: Labor Market Evidence

Overview

The decision whether and how to adjust poverty thresholds for area differences in prices depends, among other things, on the reasons that prices differ across metropolitan areas and what one believes a threshold ought to measure.[1] If one starts with the premise that a poverty threshold ought to provide the same purchasing power across markets; that is, enable one to buy the same bundle of well defined housing and non-housing goods regardless of location, then full indexing to an *appropriate* price index is in order.

If one shifts away from an emphasis on equivalent purchasing power for a fixed set of goods and services to the broader concept of equivalent individual or household well-being (utility), full indexing to area price-level differences does not follow. First, when faced with different relative prices, households adjust their consumption bundles away from those goods and services with relatively high prices and toward those with relatively low prices.[2] For a given level of well-being, household will purchase different bundles of goods and services in Peoria than in New York City.

Second, area amenities that enhance utility increase the price of land (and hence housing rents and prices), while at the same time decreasing equilibrium wages for any given level of prices. Absent full accounting for area amenities, area wages should not and do not generally increase one-for-one with area prices. That is, in a log wage equation absent control for amenities, the wage-price elasticity Θ, measured by the coefficient on lnP, is below unity. Across labor markets, real wages (W/P) and purchasing power do not equalize but, at the margin, utility does. If wages across markets were somehow administratively indexed fully with respect to P, wages would be above equilibrium in high amenity cities and below equilibrium in low amenity cities. The appropriate index for area wages (and, arguably, for adjusting poverty thresholds) is not a price index but an area wage index for workers of similar skill in jobs with similar tasks and working conditions. This is the "adjustment" that is produced more or less automatically through market forces.[3]

Similar reasoning applies to the question of whether and how to adjust poverty thresholds for price (or wage) differences. As carefully shown in Glaeser (1998), if there is sufficient mobility among the population receiving transfers (and there may not be), marginal utilities of income should equalize

[1] Throughout the discussion, I use the terms "metropolitan areas" and "cities" interchangeably. In subsequent empirical analysis, metropolitan statistical areas are identified using the Current Population Survey (CPS).
[2] The ability of households to substitute argues for area threshold adjustments using a cost of living index constructed from varying bundles of goods and services rather than using a price index based on fixed bundles.
[3] This is the logic of the (largely un-implemented) Federal area wage adjustment program that relies on an index measuring area wages by occupation within large metropolitan areas.

across cities with different price levels. Adjustment of transfers for price differences would discourage efficiency-enhancing mobility and more of taxpayers' dollars would be spent in cities where a dollar purchases fewer goods and services.

Important questions include (1) the extent to which wages for equivalent workers and jobs rise with respect to available area price indices; (2) whether recipient or poverty populations value amenities similarly to the larger population of wage and salary workers; and (3) whether low-income populations have sufficient mobility to roughly equalize marginal utilities of income (only some fraction of a population needs to be mobile to equalize marginal utilities). In this essay I provide direct evidence on question 1 and indirect evidence on question 2. I do not offer evidence on question 3, but accept the premise that low-income populations are generally less mobile than higher income populations, although greater attachment of the latter group to the labor market works in the opposite direction.

If populations near and below a poverty line have preferences with respect to area amenities that are similar to the average worker, then adjustment of poverty thresholds to a wage rather than price index would be appropriate. If populations near and below a poverty line place little or no value on area amenities, then full price adjustment of the threshold would be needed to equalize marginal utilities, but doing so would discourage desirable migration to lower cost cities. If they place greater relative value on area amenities, then it would argue for smaller partial adjustments with respect to a price index. What is most likely is that the particular bundles of public services and amenities valued by different population groups vary. Coupled with the fact that the individuals in low-income households have weak attachment to the labor market, it is difficult to generalize about the poverty population's valuation of area amenities based on evidence from the labor market.

In short, the logic from economic theory is that full adjustment of poverty thresholds to an area price index based on a fixed bundle of goods and services is inappropriate, owing both to household substitution and the valuation of amenities capitalized into prices. Full adjustment of poverty thresholds to an area wage index, however, may be appropriate. But use of an area wage index to adjust area poverty thresholds is not likely to be politically feasible. The logic of using an area wage rather than price index to adjust poverty thresholds may be persuasive only to an economist, all the more so given the weak attachment of low-income household members to the labor market.

If use of an area price index is the only feasible path to adjusting poverty thresholds it is important that the adjustment ratio approximate what would be obtained using an area wage index. As shown below, this is possible if one uses partial adjustment to prices equivalent to the wage-price elasticity across labor markets. That is, if wages for equivalent workers across cities receive wages that rise, say, 8% with respect to a 10% difference in prices (implying a wage-price elasticity of 0.80) then a poverty threshold adjustment ratio of 80% with respect to prices would, on average, approximate the use of full wage adjustment. As noted subsequently, it is essential that the same price index be used to estimate the wage-price elasticity as is used to adjust the poverty thresholds.

In what follows, I use area wage data to estimate wage-price elasticities with respect to the "CEO" price index developed by Carrillo, Early, and Olson (2010). I show that the proposed Census method for adjustment of area poverty thresholds, which accounts for housing but not non-housing costs, provides

a partial price adjustment that could roughly mimic wage indexing if it were to use the CEO index. Whether alternative area price indices, including those considered by Census, would similarly approximate wage indexation is not examined here. It is essential that before a particular price index and threshold adjustment method are adopted, one verifies that the poverty threshold adjustment ratio is highly similar to the wage-price elasticity for using that same price index. Comparing the two is relatively straightforward. Similar values imply that one approximates, on average, the poverty thresholds that would be obtained through full adjustment to an area wage index.

Census poverty threshold proposal

The Census proposal calls for experimental poverty thresholds to be fully adjusted for housing cost differences, based on the approximate 40% of household budgets spent on housing. Such adjustment would be equivalent to full price adjustment if there were no variation across cities in prices of non-housing goods. While non-housing prices vary across metropolitan areas, they do so far less than housing. Using the CEO price data (described subsequently) and a CPS sample of wage and salary workers for 264 metro areas in 2006, the coefficient of variation for the CEO price index of non-housing goods is 0.056, only a fourth as large as the 0.237 variation in the housing index.[4] So as a rough back-of-the envelope approximation, adjustment of poverty thresholds on the basis of intercity differences in housing (but not non-housing) prices results in something close (80%) to full adjustment. That is, one indexes for roughly 80% of total intercity price variation.

A virtue of the Census proposal is that it does not provide for full adjustment to prices, potentially accounting for substitution by residents in response to different relative prices and to the valuation of amenities that are capitalized in prices. Whether or not the deviation from full adjustment is the "correct" deviation can be assessed by comparing this deviation or adjustment ratio to estimates of the wage-price elasticity, assuming the same price index is used to adjust the thresholds and in estimating the elasticities. Below we provide such evidence.

Evidence on wage-price elasticities

What is the evidence on how wages vary with respect to area price differentials? I examine this using a new CEO price index developed by Carrillo, Early, and Olsen (2010), which relies on information from a large HUD Section 8 survey in 2000 that provides housing gross rental prices and highly detailed housing attributes, combined with neighborhood (census tract) information from the 2000 Decennial Census. Values for years earlier and after 2000 are calculated using BLS time series price indices for larger areas. The housing component of the price index is then combined with non-housing price information from the Council for Community and Economic Research (which produces the ACCRA index). Carrillo et al. provide estimates of a price index (including separate housing and non-housing components) for metropolitan and non-metro state areas for 1982 through 2008. I examine the relationship using the CEO price index for 2005-2008 (with pooled and separate estimates by year).

[4] Calculating average prices over the urban sample of workers is equivalent to taking the employment-weighted average across the 264 cities.

Estimated is a log wage equation, estimated at the individual worker level, of the following general form.

$$\ln W = \Theta \ln P + X\beta + \mu$$

Results for Θ, the wage-price elasticity across cities is estimated with controls for worker attributes and job sector (public/private, industry, and occupation) in order to control for worker skill and job differences across cities. I present but do not emphasize estimates of Θ based on no covariates, since such estimates do not control for systematic differences in skill across high and low price cities. Nor do I present estimates from wage specifications that include region, city size, or explicit measures of amenities, as would be appropriate if we were studying whether or not the law of one wage is approximately obtained across labor market. Exclusion of such controls is appropriate for our purposes since our goal is to see how nominal wages (and not utilities) vary with respect to price differences across cities for equivalent workers in equivalent jobs.

The analysis uses the CPS Outgoing Rotation Group (ORG) monthly earnings files for 2005 through 2008.[5] The ORG earnings supplement to the CPS includes questions on, among other things, usual earnings at the principal job the previous week, usual hours worked per week in that job, and union status. We create a measure of average hourly earnings as follows. Hourly workers report their straight-time wage rate. For hourly workers who do not report tips, overtime, or commissions, the straight time wage is used. For all other workers, the wage is measured by usual weekly earnings, which includes tips, overtime, and commissions, divided by usual hours worked per week on the principal job.[6] For workers whose weekly earnings are top-coded in the ORGs (at $2,885), we assign the estimated mean by year and gender above the cap assuming a Pareto distribution above the median.[7]

Included as controls in vector X are potential experience (in quartic form) and dummies for schooling (11), part-time, gender, marital status (2), race/ethnicity (4), foreign-born status (2), union member, public and industry sector (14), and occupation (9), and year (3).

Interest here is in estimates of Θ, the wage-price elasticity across the 264 MSAs matched in the CPS and CEO price database. About 70% of the CPS national sample resides in the 264 MSAs identified in the CPS (very small MSAs are not identified). In addition to estimates from a pooled equation with controls, separate estimates are provided by year, sex, education, and standing in the wage distribution. Corresponding estimates are provided from a specification without controls. We also compare estimates with and without inclusion of workers whose wages have been imputed in the CPS (see below).

Earnings non-response in the CPS is a serious issue and one that affects wage-price elasticity estimates. In the ORGs about 30% of surveyed individuals currently have their earnings imputed (allocated) using a cell hot deck procedure. Nonrespondents are assigned the earnings of a "similar"

[5] The CPS adopted the 2003 metropolitan area designations in May 2004, so 2005 is the first full year containing these metro definitions, while 2008 is the latest year for CEO price data.

[6] For the few workers who do not report an hourly wage and report variable hours, the wage is calculated using hours worked the previous week.

[7] Estimates compiled by Barry Hirsch and David Macpherson are posted at www.unionstats.com. Estimated means above the cap for men (women) have increased over time. In 2008 they are 1.87 (1.68) times the $2,885 cap.

respondent or donor, but match characteristics in the hot deck do not include location (not MSA, state, or region), industry, and many other important wage determinants. As shown in Hirsch and Schumacher (2004) and (Bollinger and Hirsch 2006), there is severe attenuation (referred to as "match bias") in wage equation coefficients on non-match criteria. A more complex pattern of bias exists for coefficients on imperfectly matched attributes (e.g., schooling, age, occupation). The degree of attenuation in coefficients on non-match attributes is close to the rate of earnings nonresponse. This same attenuation should apply to estimates of the wage-price elasticity Θ.

The simplest way to eliminate match bias is to omit imputed earners from the estimation sample, which is what is done here (for comparison of alternative methods, see Bollinger and Hirsch (2006)). A separate issue that arises (with or without imputed earners included) is whether there exists non-ignorable response bias. Bollinger and Hirsch (2010) address the issue of response bias using selection models. They conclude that non-ignorable response bias exists but is modest and that it does not noticeably bias slope coefficients in a wage equation. In short, inclusion of imputed earners creates serious first-order attenuation in wage equation coefficients (as seen subsequently), with or without response bias. Excluding imputations corrects for match bias. Any effects of non-ignorable response bias on estimates of Θ should be quite minor, and would not be corrected by including imputed earners.

Inclusion of imputed earners in earlier work using the CPS, coupled with use of an ACCRA price index that's more dispersed across cities than either the long discontinued BLS intercity budgets or the new CEO price index used here, leads to estimates of wage-price elasticities far below unity, on the order of .5 or less (e.g., DuMond, Hirsch, and Macpherson 1999). Winters (2009) shows that substitution of an ACS rental index for the ACCRA housing component, instrumenting the housing index to reduce measurement error, and accounting for amenities leads to wage-price elasticity estimates very close to unity (he also excludes imputed earners).

In Table 1, I show estimates of Θ, the wage-price elasticity, for the full 2005-2008 sample (minus imputed earners) for all years, plus separate estimates by gender and year. The preferred result is for the full model, which includes a relatively rich set of worker and job measures. Here the estimate of Θ is 0.86, clearly below one, but not an order of magnitude below unity as in studies using the ACCRA index and with CPS imputed earners included. An estimate of $\Theta < 1.0$ is expected given that we do not control for amenities. Finding that the wage-price elasticity is reasonably close to unity is consistent with the law of one wage and suggests that the CEO price index using rental prices is a far more appropriate measure than is ACCRA.

As seen at the bottom of Table 1, absent covariates other than dummies for year, the estimated Θ is 1.07, indicating that city wages on average rise slightly more than 1-for-1 with prices. Although informative, we should attach limited weight on results from a specification that does not control for worker skill or job sector. The higher estimate of Θ seen absent controls reflects the fact that large cities not only have higher prices, but also attract a more skilled workforce.

As a robustness check, we examine wage-price elasticity estimates by gender and year. The estimate of Θ for men is 0.835 and that for women 0.876, a difference that may be statistically significant but is not economically significant. Estimates of Θ by year are relatively stable, being 0.88,

0.85, 0.85, and 0.87 for the years 2005-2008. It will be possible to update estimates to 2009 and 2010 once CEO price data are available and to examine evidence from earlier years through a match with CPS metro definitions used prior to May 2004.

The far right column in Table 1 shows the multi-year result with imputed earners included in the CPS. The proportion of workers omitted due to imputation was 34%. This exceeds the overall rate of imputation in the CPS owing to use of the metropolitan-only sample (non-response is greater in urban areas and increases with city size). In contrast to the 0.862 wage-price elasticity obtained using the sample of CPS respondents, inclusion of imputed earners causes severe attenuation in estimates of Θ, in this case to 0.612. This 29% attenuation due to match bias is not far below the 34% non-response rate, reflecting near zero correlation of wages and prices in the sample with earnings imputed. In short, estimates of wage-price elasticities from the CPS that do not account for the match bias due to non-response will lead to severely biased (attenuated) estimates of Θ.

Properly estimated, the wage-price elasticity is much closer to unity than suggested by previous analyses, with the exception of Winters (2009) who uses a similar approach. The higher estimates result from both use of a more appropriate urban price index basing housing costs on rental equivalents and from removing the attenuation that results from imputed earnings not matched on location (among other things).

Tables 2-3 examine how wage-price elasticities vary with the level of schooling and across percentiles of the distribution using quantile regression. Schooling is a particularly rough proxy for the likelihood of poverty or near-poverty status, so greater emphasis is given to the quantile regression estimates. The evidence based on schooling level shows that estimates of Θ increase with schooling level, being 0.54, 0.80, 0.87, and 0.96 for high school dropouts, high school grads, those with some college, and those with a B.A. or above.

Quantile regression estimates of Θ, as seen in Table 3, demonstrate a similar qualitative pattern, but a narrower range of estimates. The median regression Θ is 0.84, very close to the OLS estimate of 0.86. In the tails, the 10^{th} percentile estimate of Θ is 0.69 while the estimate at the 90^{th} percentile is 1.00. Moving from the 25^{th} to 75^{th} percentiles, estimates of Θ increase only modestly, from 0.77 to 0.91. At the 33^{rd} percentile the estimate of Θ is 0.80, very close to the degree of price adjustment implicit in the Census proposal for adjusting poverty thresholds.

Absent a better understanding of why estimates of Θ are lower for those in the left tail than in the right tail of the earnings distribution, I am reluctant to attach great weight to these findings. A low wage-price elasticity in the left tail may reflect (1) that lower wage workers value more highly the non-wage amenities or other aspects of high-wage metropolitan areas such that their wages need not rise so closely with respect to prices; (2) that there is less variation in wages for lower-wage than higher wage workers owing to minimum wages or other constraints on wage adjustment ; or (3) that P is a poorer measure of relevant prices for lower than for higher wage workers and, hence, estimates of Θ are attenuated due to measurement error. Even if the wage-price elasticity estimate for lower wage workers are unbiased, we do not know how similar is the valuation of amenities among the population of *employed* wage and salary workers in the left-tail of the distribution with the valuation among the

broader and perhaps more relevant population living in poverty or near-poverty households. The latter population has low rates of labor force participation.

Taken literally, the estimates of Θ suggest that lower wage workers (and possibly the broader poverty population) place greater proportional value on area amenities than do higher wage workers. Is this plausible? It may depend on the amenities. Wages among low income households may be particularly sensitive to good public services (schools, parks, public transportation), low crime (since they may be victims), and possibly weather. Unfortunately, we know little about how the poor value the amenities that are reflected in area land prices and wages. That being said, low estimates of Θ for the low-wage population reinforce the previous conclusion that if poverty thresholds are to be adjusted by prices, partial rather than full adjustment is in order.

Implications and Conclusion

Full adjustment of poverty thresholds with respect to an area price index would be highly problematic if the wage-price elasticity with respect to the designated index were substantially below unity. Wages for workers of a given skill do not adjust fully to area price differences owing to amenities (which raise prices and lower real wages) and because consumers have the ability to vary their consumption bundles in response to differences in relative prices. Previous evidence using the ACCRA price index and estimates in which imputed earners are included in the CPS produced elasticity estimates substantially below unity, on the order of one half. As seen in Winters (2009) and shown here using the CEO price index, wage-price elasticities using the CPS are only moderately lower than unity, on the order of 0.80. Elasticities of this magnitude (i.e., not far below unity) are obtained when the housing cost component of the price index is based on rental prices and when imputed earners are removed from the analysis.

An economic case can be made for full adjustment of poverty thresholds using an area *wage* index. Use of a wage rather than price index to adjust poverty thresholds, however, would not be readily understandable to the public or to most policy makers. Partial adjustment of poverty thresholds using an area price index, however, can roughly mimic wage indexing if the adjustment ratio for the threshold is similar to the wage-price elasticity. Although there is variation across cities in how their wages rise with respect to prices, raising poverty thresholds with respect to prices by a partial adjustment ratio close to Θ will get things right, at least on average. In determining whether or not the adjustment ratio and the wage-price elasticity Θ are similar, it is essential that the same price index be used for the poverty adjustment and for estimates of Θ.

The Census proposal to adjust poverty thresholds fully with respect to housing price differences but not with respect to non-housing prices appears to result in a partial adjustment factor on the order of 75%-80%, based on my rough estimate. This matches estimates of the wage-price elasticity that I obtain using CPS wage data and the CEO area price index. Thus, Census adjustment of poverty thresholds using the CEO housing index would provide an adjustment that on average would approximate full adjustment by a wage index. On this basis, I would support the Census adjustment method were it to use the CEO price index. Apart from the issues discussed in this essay, a case can be made that the CEO area price index provides substantial advantages as compared to alternative indices (Carrillo, Early, and Olsen,

2010). The Census adjustment method should not be used with alternative area price indices (including those proposed by Census) unless it can be shown that wage-price elasticities using these price indices are similar to the threshold adjustment ratio.

References

Bollinger, C.R. and B.T. Hirsch. 2006. "Match Bias from Earnings Imputation in the Current Population Survey: The Case of Imperfect Matching." *Journal of Labor Economics* 24(3): 483-519.

Bollinger, C.R. and B.T. Hirsch. 2010. "Is Earnings Nonresponse Ignorable?" IZA Discussion Paper 5346, November.

Carrillo, P.E., Early, D.W., E.O. Olsen. 2010. "A Panel of Price Indices for Housing, Other Goods, and All Goods for All Areas in the United States 1982-2008," working paper.

DuMond, J.M., B.T. Hirsch, D.A. Macpherson. 1999. "Wage Differentials Across Labor Markets and Workers: Does Cost of Living Matter?" *Economic Inquiry* 37(4), October, 577-98.

Glaeser, E.L. 1998. "Should transfer payments be indexed to local price levels?" *Regional Science and Urban Economics* 28(1): 1-20.

Hirsch, B.T. and E.J. Schumacher. 2004. "Match Bias in Wage Gap Estimates Due to Earnings Imputation," *Journal of Labor Economics* 22(3): 689-722.

Winters, J.V. 2009. "Wages and prices: Are workers fully compensated for cost of living differences? 2009. *Regional Science and Urban Economics* 39: 632-643.

Table 1: Wage-Price Elasticity Estimates from the CPS

	All	Men	Women	2005	2006	2007	2008	with imputations
Full model:								
Θ	**0.862**	**0.835**	**0.876**	**0.878**	**0.846**	**0.854**	**0.871**	**0.612**
s.e.	0.067	0.064	0.073	0.071	0.077	0.064	0.064	0.051
Model with year dummies, no controls:								
Θ	1.067	1.029	1.104	1.074	1.046	1.055	1.093	0.783
s.e.	0.130	0.145	0.128	0.131	0.139	0.136	0.122	0.097
N	338,846	170,900	167,946	84,445	85,047	84,968	84,386	513,482

Data: 2005-2008 CPS-ORG files and the CEO Price Panel. The dependent variable is the log of hourly earnings. Shown are estimates of theta, the coefficient on lnP, measured for each of 264 MSAs by year. The "Full models" include potential experience (in quartic form) and dummies for schooling (11), part-time, gender, marital status (2), race/ethnicity (4), foreign-born status (2), union member, public and industry sector (14), and occupation (9), and year (3). Standard errors are clustered on MSA.

Table 2: **Wage-Price Elasticity Estimates by Education Group**

	All	Dropouts	High School	Some College	College+
Full model:					
Θ	**0.862**	**0.555**	**0.805**	**0.874**	**0.964**
s.e.	0.067	0.094	0.060	0.075	0.093
Model with year dummies, no controls:					
Θ	1.067	0.447	0.671	0.793	1.059
s.e.	0.130	0.086	0.074	0.077	0.100
N	338,846	36,248	87,960	98,287	116,351

See note to Table 1.

Table 3: **Quantile Regression Estimates of Wage-Price Elasticities**

	Percentile						
	10^{th}	25^{th}	33^{rd}	Median	67^{th}	75^{th}	90^{th}
Full model:							
Θ	**0.686**	**0.773**	**0.803**	**0.843**	**0.874**	**0.906**	**0.999**
s.e.	0.010	0.008	0.008	0.008	0.008	0.009	0.013
Model with year dummies, no controls:							
Θ	0.487	0.771	0.951	1.115	1.278	1.379	1.473
s.e.	0.012	0.012	0.013	0.014	0.013	0.014	0.017

See note to Table 1. Standard errors are not clustered. Sample size is 338,846.

Report on Adjusting Poverty Thresholds for Geographic Price Differences

Edgar O. Olsen

University of Virginia

April 21, 2011

Revised May 9, 2011

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 criteeria), 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.

[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.

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.

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 housing prices vary much less across areas than the prices of other goods, 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.

Comments on the U.S. Census Bureau and Interagency Technical Working Group Supplemental Poverty Measure

by

Stuart S. Rosenthal
Maxwell Advisory Board Professor of Economics
426 Eggers Hall, Syracuse University
Syracuse, NY 13244-1020
Phone: (315) 443-3809; Fax: (315) 443-1081
Email: ssrosent@maxwell.syr.edu
Home Page: http://faculty.maxwell.syr.edu/ROSENTHAL/

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

Revised May 9, 2011

1. Overview

Longstanding government measures of poverty are based on income thresholds that are uniform across locations. The newly developed supplemental poverty measure (SPM) departs from that tradition by setting higher income thresholds in metropolitan and rural areas with higher housing costs. This allows for inter-metropolitan differences in housing prices that affect the level of income necessary for a family to purchase a minimum level of food, clothing, shelter, and other essential items necessary to attain a minimum standard of living. Numerous government programs that provide assistance to the poor define eligibility in part based on a family's designated poverty status. Although the SPM is not being used for such purposes, the possibility exists that it may be used in this fashion at some point in the future. In that regard, the SPM has potential to affect not only perceptions of regional differences in poverty, but also the geographic distribution of benefits from federal and state government programs that provide assistance to the poor.

Central to the new SPM measure of poverty is the idea that the level of income that defines a given size family's poverty status should be higher in metropolitan (and non-metropolitan) areas with high housing costs. For these purposes, housing costs are defined so as to include expenditures on utilities including electricity, heating, etc.. This is presumably motivated by recognition that the need for heating/cooling as well as energy prices differ widely across geographic areas as between Houston versus Detroit, for example. The new poverty measure also measures housing costs differently for renters, owner-occupiers with a mortgage, and owner-occupiers without a mortgage. This feature of the SPM is motivated by observable differences in the annual out-of-pocket costs of housing that vary with each of the three housing tenures just noted. Calculation of the SPM, at least in its current form (e.g. Renwick (2011)), is based on either mean or median housing costs for individual identified locations.

My comments below are organized into three subsections. I will first offer some perspectives intended to critique the SPM. This is followed by one suggestion for an alternative feasible method of calculating the SPM (to use hedonic-based measures of quality adjusted housing costs). Suggestions for further research on geographic adjustment of poverty thresholds are provided in the final section, as are recommendations for the SPM going forward.

2. Critique of the SPM

I will comment here on several issues that affect interpretation of the SPM and related measures. First among these is what drives spatial variation in housing costs and whether local residents obtain something valuable in exchange for living in an expensive area. This is relevant for spatial variation in housing costs across metropolitan areas and also within individual cities. This question also goes to the very core of the SPM and has implications for the geographic scale over which variation in housing costs should be taken into account when setting poverty income thresholds. A second set of issues concerns housing tenure status. Families choose whether to live in owner-occupied housing, and if so, the size of the mortgage to hold. Here too there are questions about what drives a low-income family's choice of housing tenure and whether the family might obtain something valuable in exchange for choosing a seemingly more expensive way of meeting their housing needs. Each of these sets of issues are considered below.

2.1 Do families derive benefits in return for living in more expensive locations?

At the outset, it is important to recognize that housing values and rents are driven by land values and construction costs. The costs of construction include materials and labor and are quite similar across locations, at least in comparison to lot values. This suggests that raw land and housing values vary spatially primarily in response to the presence of local attributes that are typically not traded directly on

the market. Such attributes include everything from air quality to school quality, scenic views, access to public transportation, and perhaps most importantly, proximity to employment. These location-specific specific attributes are valued by households and firms in differing degrees contributing to spatial variation in housing costs. Two questions immediately arise in regards to the SPM: (i) to what extent do families derive benefits from living in expensive areas and (ii) should such benefits and related differences in housing costs be taken into account when measuring poverty?

It is also important to recognize that adjusting poverty measures for spatial differences in the cost of housing changes the poor's relative cost of access to attributes in expensive metro areas relative to inexpensive locations. That is because some state and federal assistance programs for the poor are dependent on a family's designated poverty status. Increasing the poverty income threshold in expensive areas has the potential to increase welfare benefits in expensive cities relative to inexpensive locations. To the extent this occurs that may attract or help to retain low-income families in high-cost locations. It is desirable, therefore, to consider the extent to which the SPM might induce shifts in the location of the poor, and possible unintended consequences from such behavior.

In both instances – whether the poor derive benefits from living in a more expensive location, and whether spatially-based poverty measures might attract the poor to high-cost areas – a key underlying consideration is the extent to which the poor are mobile and choose their locations. In part, this is a question of time horizon since over a sufficiently long horizon most families are mobile, at least to some degree. This is also a question of relocation costs and access to alternative locations. Moving costs include the possible loss of valuable locally-based social networks including family and friends that may be concentrated in a particular metropolitan area. For many minorities including especially recent immigrants, this could also include the presence of ethnic/racial enclaves in which social and economic support networks may be more accessible. Other costs include access to car travel which frees up the auto owner from reliance on foot transport and public transit. This latter issue is of special importance for low-income families as the very poor are far less likely to own a car than are higher income families.

On this latter point, consider Table 1 below which is taken from Rosenthal (2008a). The table presents regressions in which the dependent variable is the percent of occupied housing units in a census tract that own a car. Control measures include distance to the central business district (CBD) which proxies for the local density of development and also the census tract poverty rate. Estimates in the first three columns are based on data from the 1980, 1990, and 2000 censuses, respectively. These regressions also control for MSA fixed effects. The fourth column pools data across the three decades and includes MSA by year fixed effects.

Notice that the coefficients on tract poverty rate are always close to -1.0 for each of the regressions. This indicates that as an approximation, a tract populated entirely with families below the poverty line would have close to zero car ownership – consistent with the idea that families in poverty are far less likely to own cars. These estimates confirm that very low income families in urban areas are generally reliant on walking and public transit for travel. It is worth emphasizing that public transit is more effective and common in densely developed areas. Glaeser, Kahn, and Rappaport (2007) draw on both of these stylized facts in arguing that access to public transit helps to explain why the urban poor are disproportionately concentrated in the central cities.

Two sets of models in the urban economics literature suggest that provided households (and firms) are mobile, spatial variation in housing costs reflect differences in local attributes including especially differences in proximity to employment opportunities. The first of these are within-city models developed by Mills and others in the late 1960s and early 1970s. These models tend to ignore within-city variation in wages which is typically thought to be relatively modest. A core point in these models is that house prices fall with distance to employment centers so as to compensate workers for longer, more

costly commutes. More generally, house prices fall with distance to *any* attractive attribute, of which proximity to employment is just one such example. Other valued local attributes include things like proximity to public transit, school quality, safe neighborhoods, scenic views, etc. as noted above.

Table 1: Car Ownership and Poverty
(Dependent Variable: Percent of census tract occupied housing units that own a car)[a]

	Individual Decades			Balanced Panel
	1980	**1990**	**2000**	**1980-2000**
Tract Poverty Rate	-1.0432	-0.8692	-0.8001	-0.8867
	(-300.36)	(-301.42)	(-265.77)	(-265.07)
Distance to CBD	0.0022	0.0019	0.0015	0.0019
	(61.51)	(57.49)	(47.62)	(75.35)
Constant	0.9711	0.9647	0.9657	0.9679
	(1,340.51)	(1,396.17)	(1,374.65	(1,855.94)
Observations[b]	48,950	50,312	50,511	145,590
Adj. R-squared	0.7573	0.7510	0.7242	0.7380
MSA Fixed Effects[c]	325	331	331	-
MSA*year FE[b]	-	-	-	975

[a]t-ratios are based on robust standard errors.
[b]Samples are restricted to census tracts in MSAs.
[c]Fixed effect values reflect deviations from the sample-wide constant.

Source: Rosenthal (2008a).

A different set of models originates with Roback (1982) and subsequent empirical applications by Blomquist et al (1989) and Gyourko and Tracy (1991). These papers consider inter-city differences in housing costs and wages. In this literature – sometimes referred to as the quality of life literature – cross-city differences in attributes are capitalized into metro-level house prices and wages. This arises because inter-city differences in attributes shift the supply and demand for labor across areas as workers and firms seek more attractive/profitable locations. From the worker's perspective, real wages adjust across cities to ensure a spatial equilibrium that eliminates incentives for similar-type workers to relocate. In previous work (Gabriel and Rosenthal (2005) and Chen and Rosenthal (2008)), I further emphasize that the sum of labor and land costs will also adjust across cities to ensure that firms do not have any incentive to relocate. In these models, lower real wages signal the presence of local amenities that are valued by the workers. Analogously, high location-specific input costs (as measured by the sum of labor and land costs) signal the presence of local attributes that are valued by firms. Because households choose locations to maximize wellbeing while firms choose locations to maximize profits, households and firms need not favor the same locations. Results from my previous work in this area (Gabriel and Rosenthal (2005) and Chen and Rosenthal (2008)) confirm this: households often favor coastal cities in relatively temperate climates and also some rural areas, while firms tend to favor large metropolitan areas with agglomerations of existing employment and industry (e.g. Rosenthal and Strange (2004)).

Taken at face value, an important segment of the prior urban literature could be interpreted as to suggest that there may be little basis for adjusting measures of poverty for differences in housing costs across metropolitan areas. Instead, differences in housing costs across locations reflect differences in valued local attributes that are implicitly "purchased" by households through their choice of where to live. There is, however, an important "but." If poor households face constraints that restrict their mobility then it is possible that market prices (including both wages and housing costs) would not adjust to compensate the

poor for differences in attributes across locations. An example of an entire literature predicated on this idea is the "spatial mismatch" literature (e.g. Holzer (1991), Kain (1968, 1992)).

Early versions of the spatial mismatch literature argued that suburbanization of jobs coupled with suburban housing market discrimination against blacks impeded African American employment opportunities. Although more recent portions of the literature have broadened in focus, the basic structure remains: constraints on mobility prevent market prices from fully compensating lower-income (and minority) individuals for differences in proximity to different local attributes, including proximity to employment. In a paper with Gabriel (Gabriel and Rosenthal,1998), I explore this question using data from the 1980s. Results indicate that commute distance to work is longer for urban blacks as compared to comparable white workers. That difference is diminished when neighborhood-level attributes and housing costs are taken into account (through the use of neighborhood fixed effects) but a race-related difference remains.

Consistent with the discussion above, implicit in the SPM is the assumption that very low-income families are relatively immobile and/or face constraints on their residential opportunities. Those constraints prevent poor families from being fully compensated for differences in local attributes through spatial variation in housing costs (and wages). To the extent this is true, then low-income families living in expensive locations will only be partly compensated for their higher housing costs through access to valued local attributes (e.g. better employment opportunities, higher quality schools, etc.). Under such conditions it would be desirable to adjust poverty income thresholds for local differences in housing costs.

This discussion points to two questions that bear on interpretation and implementation of the SPM. The first is whether the poor face notable constraints that limit their residential opportunities. The second is whether the set of local attributes that are capitalized into local house prices are ones that should be taken into account when considering whether a family is classified as not being able to afford basic "necessities." For example, should proximity to employment and shorter commutes be treated as a necessary "commodity"? Should a scenic view be considered, or proximity to better schools?

The spatial mismatch literature highlights the role of racial/ethnic segregation and discrimination as one possible source of such restrictions that might limit the residential opportunities of the poor. Another source could be local zoning laws that mandate large lots and low-density development. A very different type of restriction could be the limited ability of the poor to afford a vehicle as documented above. Especially in urban areas this would tend to push the poor to inner city location in order to gain better access to public transit (e.g. Glaeser, Kahn, and Rappaport (2007)). Yet another possible constraint on housing opportunities for the poor is the age of the local housing stock. Higher income families tend to occupy newer homes while the poor tend to occupy old homes, at least if they are not in subsidized housing. Because cities develop from the center outwards over time, and redevelop from the center outwards over time, this affects the age of the housing stock at different locations, and therefore, the location of housing opportunities for low-income families (e.g. Brueckner and Rosenthal (2009), Rosenthal (2008b)). If any of these factors impede the ability of low-income families to consider alternative locations within or across cities, then the possibility exists that the poor may live on "islands" that are relatively disjoint from the rest of the housing market. In that event, it is possible that local attributes may not fully compensate lower income families for differences in house prices and wage within and across cities.

A final constraint of a completely different nature should also be mentioned. The very poor tend to have limited labor market skills and are likely to work at minimum wage jobs when employed. Intercity quality of life models, however, are based on the idea that *both* house prices and nominal wage adjust to offset between-city differences in amenities. This would not occur for minimum wage workers. Instead,

even if the poor are mobile, differences in real wages across metropolitan areas would be driven primarily by differences in housing costs. Although it is possible that housing price (rent) capitalization could still fully offset cross-city differences in amenities, minimum wage constraints may impede such outcomes. This is something that requires further consideration as it provides an additional possible argument for taking cross-city variation in housing costs into account when setting poverty income thresholds.

2.2 Do family's derive benefits from choosing a more "expensive" form of housing tenure?

As noted earlier, the SPM calculates different poverty income thresholds for renters, owner-occupiers with a mortgage, and owner-occupiers without a mortgage. This section offers some comments and perspective on that feature of the SPM.

It is important to recognize that housing tenure is for many families a choice. This is almost certainly the case for low-income homeowners who could otherwise rent. It is doubly the case for many low-income homeowners without a mortgage as many of these families could likely have secured some form of a mortgage or could have also choosen to rent. Among low-income homeowners with a mortgage there is also often discretion as to the size of the mortgage, loan term (e.g. 15 year, 30 year, etc), and loan type (e.g. fixed, adjustable). These features affect the size of the monthly mortgage payments.

The table below offers some perspective by reporting rental rates among families below the poverty line, disaggregated by urban/rural locations (see Rosenthal (2008a) for details). From the table, it is noteworthy that among the urban poor, rental rates are quite high, at roughly 81 percent. Nevertheless, that still leaves 19 percent of the urban poor in owner-occupied housing. For families in less densely developed areas rental rates among the poor are far lower: roughly 64 percent in non-central city urban areas and 55 percent in non-metropolitan areas. These numbers indicate that many families living in poverty own their own homes.

Table 2: Rental Rates Among Families Living in Poverty in the United States[a]

Entire U.S.	Not In Metro Area	In Metro Area Outside Central City	In Metro Area Central City
68.67%	55.25%	63.88%	80.82%

[a]Sample excludes families in mobile homes. Estimates are based on household-level data from the 2000 Decennial Census. Household weights are used to ensure results are representative of the entire U.S.

Given the mix of housing tenure among the poor, on the surface it would seem natural to allow for differences in housing costs dependent on whether a family rents, owns without a mortgage, or owns with a mortgage. Nevertheless, housing tenure including the size of the mortgage is a choice and this has possible implications for interpretation of SPM poverty income thresholds. Previous studies, for example, have documented that most homeowners take on more mortgage debt than is needed to purchase their home given the family's wealth relative to home value. Findings from this literature indicate that a portion of the "excess" mortgage debt is used for consumption and home improvements, but an important share is also used for investments of various sorts. See, for example, Jones (1994) and Canner et al (2002).

Although it is tempting to view the poor as having no wealth and therefore entirely reliant on financing to purchase a home, this is likely an oversimplification for some households. The SPM recognizes this by

treating homeowners without a mortgage differently from homeowners with a mortgage. An open and relevant question is the extent to which the poor might hold discretionary amounts of mortgage debt. It is also important to know what sorts of non-housing expenditures and/or investments the poor might use that additional debt for should they hold larger mortgages than are necessary to finance their homes. It is also important to recognize that homeowners essentially rent to themselves, and thereby receive a flow of imputed income from the home.

To illustrate, consider the following hypothetical example. Suppose three low-income families live in identical adjacent homes. The homes are completely durable and require no maintenance. There are no property or income taxes. All three families have zero earned income. One family rents, one owns without a mortgage, and one owns with a 100 percent interest-only mortgage. Suppose that each of the homes is valued at $100,000 and each family has a net worth of $100,000. Mortgage rates are 3 percent as are rates on U.S. treasury bonds. Market rents are $3,000 (equal to the discount rate multiplied by house value). The family that rents invests their wealth in treasury bonds and earns $3,000 interest each year, just enough to cover their housing rent. The family that owns without a mortgage receives no interest income and has no out of pocket housing expenses: this family implicitly rents to itself and is equally well off financially as the renter. The family that owns with a 100 percent mortgage pays $3,000 in mortgage interest each year. The family invests both the $100,000 mortgage and its original $100,000 of net wealth in treasury bonds. This yields a total investment of $200,000 which generates $6,000 interest earnings. After paying the $3,000 interest owed to the mortgage lender, this family enjoys a net investment income of $3,000. This family is therefore better off financially because it leveraged its home purchase and invested the mortgage in a risk free asset.

The example above is very stylized but illustrates two important points. First, owner-occupiers that own their homes without a mortgage effectively rent to themselves. Absent taxes or other considerations, such families face the same net housing costs as renters. Second, for owner-occupiers with discretionary levels of wealth relative to the value of their home, there is a financial incentive to lever up on the mortgage. This is likely especially relevant for families that hold 2nd mortgages, including home equity lines of credit and related loan instruments. The return from levering up is also amplified if the home appreciates in value (a scenario not shown here). Although on the surface, holding a larger mortgage may appear to increase a household's housing costs, the illustration above suggests that higher monthly mortgage payments could be offset or even more than compensated through an increase in investment earnings or other uses of the mortgage debt that are highly valued by the family. The SPM, however, sets different poverty income thresholds for renters and owners, including *higher* income thresholds for homeowners with a mortgage. These issues should receive further attention.

A final perspective on the tenure issues above is that the decisions to own a home and the chosen level of mortgage debt are both tied to a family's level of wealth. Wealth is not observed in the ACS. Although the ACS does provide some information on the family's mortgage, details on the mortgage are limited. This suggests that it will be difficult to adequately address issues related to size or mortgage when using the ACS to calculate the SPM.

3. Measuring Housing Costs in the SPM

The current plan for the SPM as proposed by the ITWG measures metropolitan level and rural housing costs using median or mean rents and prices for the different locations (see Renwick (2011), page 9). Although such an approach will likely yield informative comparisons, a better option seems to be available when working with the ACS. Specifically, because the ACS is a very large sample, it should be possible to run hedonic regressions of log rent and log price controlling for housing attributes that are reported in the survey in addition to location fixed effects. The location fixed effects would then capture

regional differences in *quality adjusted* house prices and rents. For examples of this approach, see previous work of mine on quality-of-life measures with Gabriel and Chen (Gabriel and Rosenthal (2005), Chen and Rosenthal (2008)).

In the ACS, one can identify the following attributes of the homes (see www.ipums.org for details):

- presence of kitchen or cooking facilities
- presence of a refrigerator
- sink with faucet
- presence of a stove
- number of rooms
- complete plumbing facilities
- number of rooms
- presence of hot and/or cold piped in water
- presence of a bathtub or shower
- presence of a flush toilet
- age of the structure
- number of housing units in the structure
- number of bedrooms

Although this is clearly an incomplete list of the attributes of any given home, it may well be sufficient to explain much of the variation in the quality of the structure that a family occupies, and especially when considering lower income housing. The location fixed effect can then be used to highlight the degree to which quality adjusted prices and rents vary across locations. To see how, consider the following.

One would first run regressions of the following sort:

$$\log(r_{s,jt}) = \gamma_{jt}^{r} + \theta_{t}^{r} Z_{s}^{r} + u_{s,jt}^{r} \tag{3.1}$$

where r is the rent for house s in location j and time period t. Z is the vector of housing attributes described above while γ_{jt}^{r} is the location fixed effect. The regression would be run separately for each time period allowing γ_{jt}^{r} to vary across sample years.

One would next need to form a fixed bundle of housing attributes that would be used to specify a "standard" type home, $\bar{Z}^{r}$. The rent (price) for a standard home in location j in period t is then formed by specifying:

$$\bar{r}_{jt} = \exp[\gamma_{jt}^{r} + \theta_{t}^{r} \bar{Z}^{r}] \qquad . \tag{3.2}$$

Quality adjusted rents (prices) based on (3.2) could then be augmented with further measures of utility costs as described in Renwick (2011).

Controlling for house quality in this fashion would be preferable to using median or mean housing costs, and should be feasible given the large sample sizes in the ACS. A reason not to do this would be if the type of housing stock available for the poor to occupy differed widely across metropolitan and rural areas. At least for urban areas, this seems like a relatively minor concern. However, for urban versus rural

comparisons, there could be notable differences in access to mobile homes versus multi-family housing, so this is something worth considering further.

4. Areas for Further Research and Recommendations

4.1 Future research

Two areas for further research really stand out. The first is research on the mobility of the poor. How frequently do very low income families move and over what geographic distance? When the poor move do they go to a different neighborhood within the same city, or to a different metropolitan area altogether? As a general principle, differences in housing costs across locations should be taken into account when setting poverty income thresholds provided one of two conditions are met: (i) if low-income families are not fully compensated for higher cost locations through improved access to valued local attributes *or* (ii) local attributes that prompt higher housing costs are not of a nature that lift low-income families out of poverty (e.g. scenic views). From this perspective, controlling for location at the level of a city block would clearly be too refined a geographic level. That is because mobility costs are likely low at the block level and differences in housing costs between city blocks reflect differences in block-level amenities. An open question, however, is over what distance or geographic unit (e.g. county, metro areas state) are low-income family residential locations approximately exogenously and not subject to choice? A further challenging issue is to specify those local attributes that are not typically considered necessary to help lift a low-income family out of poverty, as with scenic views and proximity to a beach, for example.

A second area of further research is to assess the degree to which the poor hold discretionary amounts of mortgage debt and use that additional debt to finance other non-housing activities, including possibly non-housing investments. Although the stereotypic image of a family living in poverty is one that has little income and also little wealth, some low-income families may nevertheless hold notable stores of wealth. An example would be house-rich cash-poor elderly. For low-income families with some wealth, the size of the mortgage may be an important choice that has implications for whether the family is truly forced by circumstances beyond their control to live in poverty.

A third area worth investigating is the extent to which the poor migrate to cities that offer more generous social safety nets. This has implications for possible use of the SPM to govern eligibility for state and federal programs designed to provide assistance to the poor.

A fourth place for further research is to explore the use of regression-based methods for measuring quality adjusted housing rents and prices across locations using the ACS.

4.2 Recommendations

Should the SPM adjust poverty income thresholds for cross-location differences in housing costs? I believe that the answer is yes because constraints on mobility likely do prevent the poor from being fully compensated for cross-metropolitan differences in housing costs. However, the extent of adjustment is an open question and almost surely lies between the two extremes of zero adjustment (as with the current official poverty measure) and 100 percent adjustment (as with the current proposed version of the SPM). Instead, the level of adjustment should depend on the extent to which the poor are *forced* – because of mobility constraints – to implicitly pay for local attributes that policy makers deem unnecessary for the poor to attain a minimum level standard of living, as with a scenic view, for example. If instead the poor *choose* to locate in a more expensive area that provides valuable services, such as proximity to public transit, then less adjustment for poverty income thresholds would seem appropriate.

In addition, while a case can be made for setting poverty income thresholds at different levels for renters, owners without a mortgage, and owners with a mortgage, I believe the case is fairly weak and difficult to implement. Instead, a more expedient and defensible approach is to measure housing costs as if all low-income families are renters. This approach would implicitly recognize that the decision to own a home is typically a choice, as is often the size and type of mortgage.

References

Blomquist, Glenn, Mark Berger, and John Hoehn (1988), "New Estimates of the Quality of Life in Urban Areas," *American Economic Review*, 78, 89-107.

Brueckner, Jan and Stuart Rosenthal (2009). "Gentrification and Neighborhood Cycles: Will America's Future Downtowns Be Rich?" *Review of Economics and Statistics*, 91(4), 725-743.

Canner, Glenn, Karen Dynan, and Wayne Passmore (December 2002), "Mortgage Refinancing in 2001 and Early 2002," *Federal Reserve Bulletin*, 469-481.

Chen, Yong and Stuart Rosenthal (2008). "Local Amenities and Life Cycle Migration: Do People Move for Jobs or Fun?" *Journal of Urban Economics*, 65(3), 519-537.

Gabriel, Stuart and Stuart Rosenthal (1996). "Commute Times, Neighborhood Effects, and Earnings: An Analysis of Compensating Differentials and Racial Discrimination," *Journal of Urban Economics*, 40: 61-83.

Gabriel, Stuart A. and Stuart S. Rosenthal (2004). "Quality of the Business Environment Versus Quality of Life: Do Firms and Households Like the Same Cities?" *The Review of Economics and Statistics*, 86(1): 438-444 (February).

Glaeser, Edward L., Mathew Kahn, and Jordan Rappaport (2008). Why Do the Poor Live in Cities? The Role of Public Transportation, *Journal of Urban Economics*, 63, 1-24.

Gyourko, Joseph and Joseph Tracy (1991), "The Structure of Local Public Finance and the Quality of Life," *Journal of Political Economy*, vol. 99, no. 4, 774-806.

Holzer, Harry (1991). "The Spatial Mismatch Hypothesis: What Has the Evidence Shown." *Urban Studies*, 28(1), 105-122.

Jones, Larry (1994), Home Mortgage Debt Financing of Nonhousing Investments, *Journal of Real Esatte Finance and Economics*, 9, 91-112.

Kain, John F. (1968), "Housing Segregation, Negro Employment, and Metropolitan Decentralization," *The Quarterly Journal of Economics*, 82(2), 175-197.

Kain, John F. (1992), "The Spatial Mismatch Hypothesis: Three Decades Later," *Housing Policy Debate*, 3(2), 371-392.

Renwick, Trudi (2011), "Geographic Adjustments of Supplemental Poverty Measure Thresholds: Using the American Community Survey Five-Year Data on Housing Costs," U.S. Census Bureau *mimeo*.

Roback, Jennifer (1982), "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 90, 1257-78.

Rosenthal, Stuart S. and William Strange (2004). "Evidence on the Nature and Sources of Agglomeration Economies", *in* the Handbook of Urban and Regional Economics, Volume 4, pg. 2119-2172, Elsevier, eds. Vernon Henderson and Jacques Thisse.

Rosenthal, Stuart (2008a). “Where Poor Renters Live in Our Cities: Dynamics and Determinants,” *in* “Revisiting Rental Housing: Policies, Programs, and Priorities,” Nicolas Retsinas and Eric Belsky, *eds*., Brookings Press, 59-92.

Rosenthal, Stuart (2008b). “Old Homes, Externalities, and Poor Neighborhoods: A Model of Urban Decline and Renewal,” *Journal of Urban Economics*, 63(3), 816-840.

Appendix 4: Keynote Presentation for Cost of Living and the Supplemental Poverty Measure

Angus Deaton

PRICE INDICES AND THE MEASUREMENT OF WELL-BEING

Why am I here?

- I have not worked on geographical price differences and poverty in the US
- However:
  - I was a member of the NAS panel on poverty and family assistance
  - I was a member of an NAS panel on cost-of-living and price indexes
  - I have been involved for many years in the measurement of global poverty
    - ICP price index work that supports it
- Bring some of that experience to a broad discussion of today's topics

NAS panel versus now

- Immense amount of high quality work done on the geographical pricing issue
- The NAS panel did not give this very deep consideration
  - It just seemed obvious that it needed to be done
  - We should have read and thought more!
  - Since then, predominant view has been
    - Wide agreement on need to make the adjustment
    - But no one knew how to do it (GAO report)
  - Now we know how to do it
    - So the question is whether we *want* to do it
- It is clearly *important*: moving poverty towards the coasts involves very large numbers of people

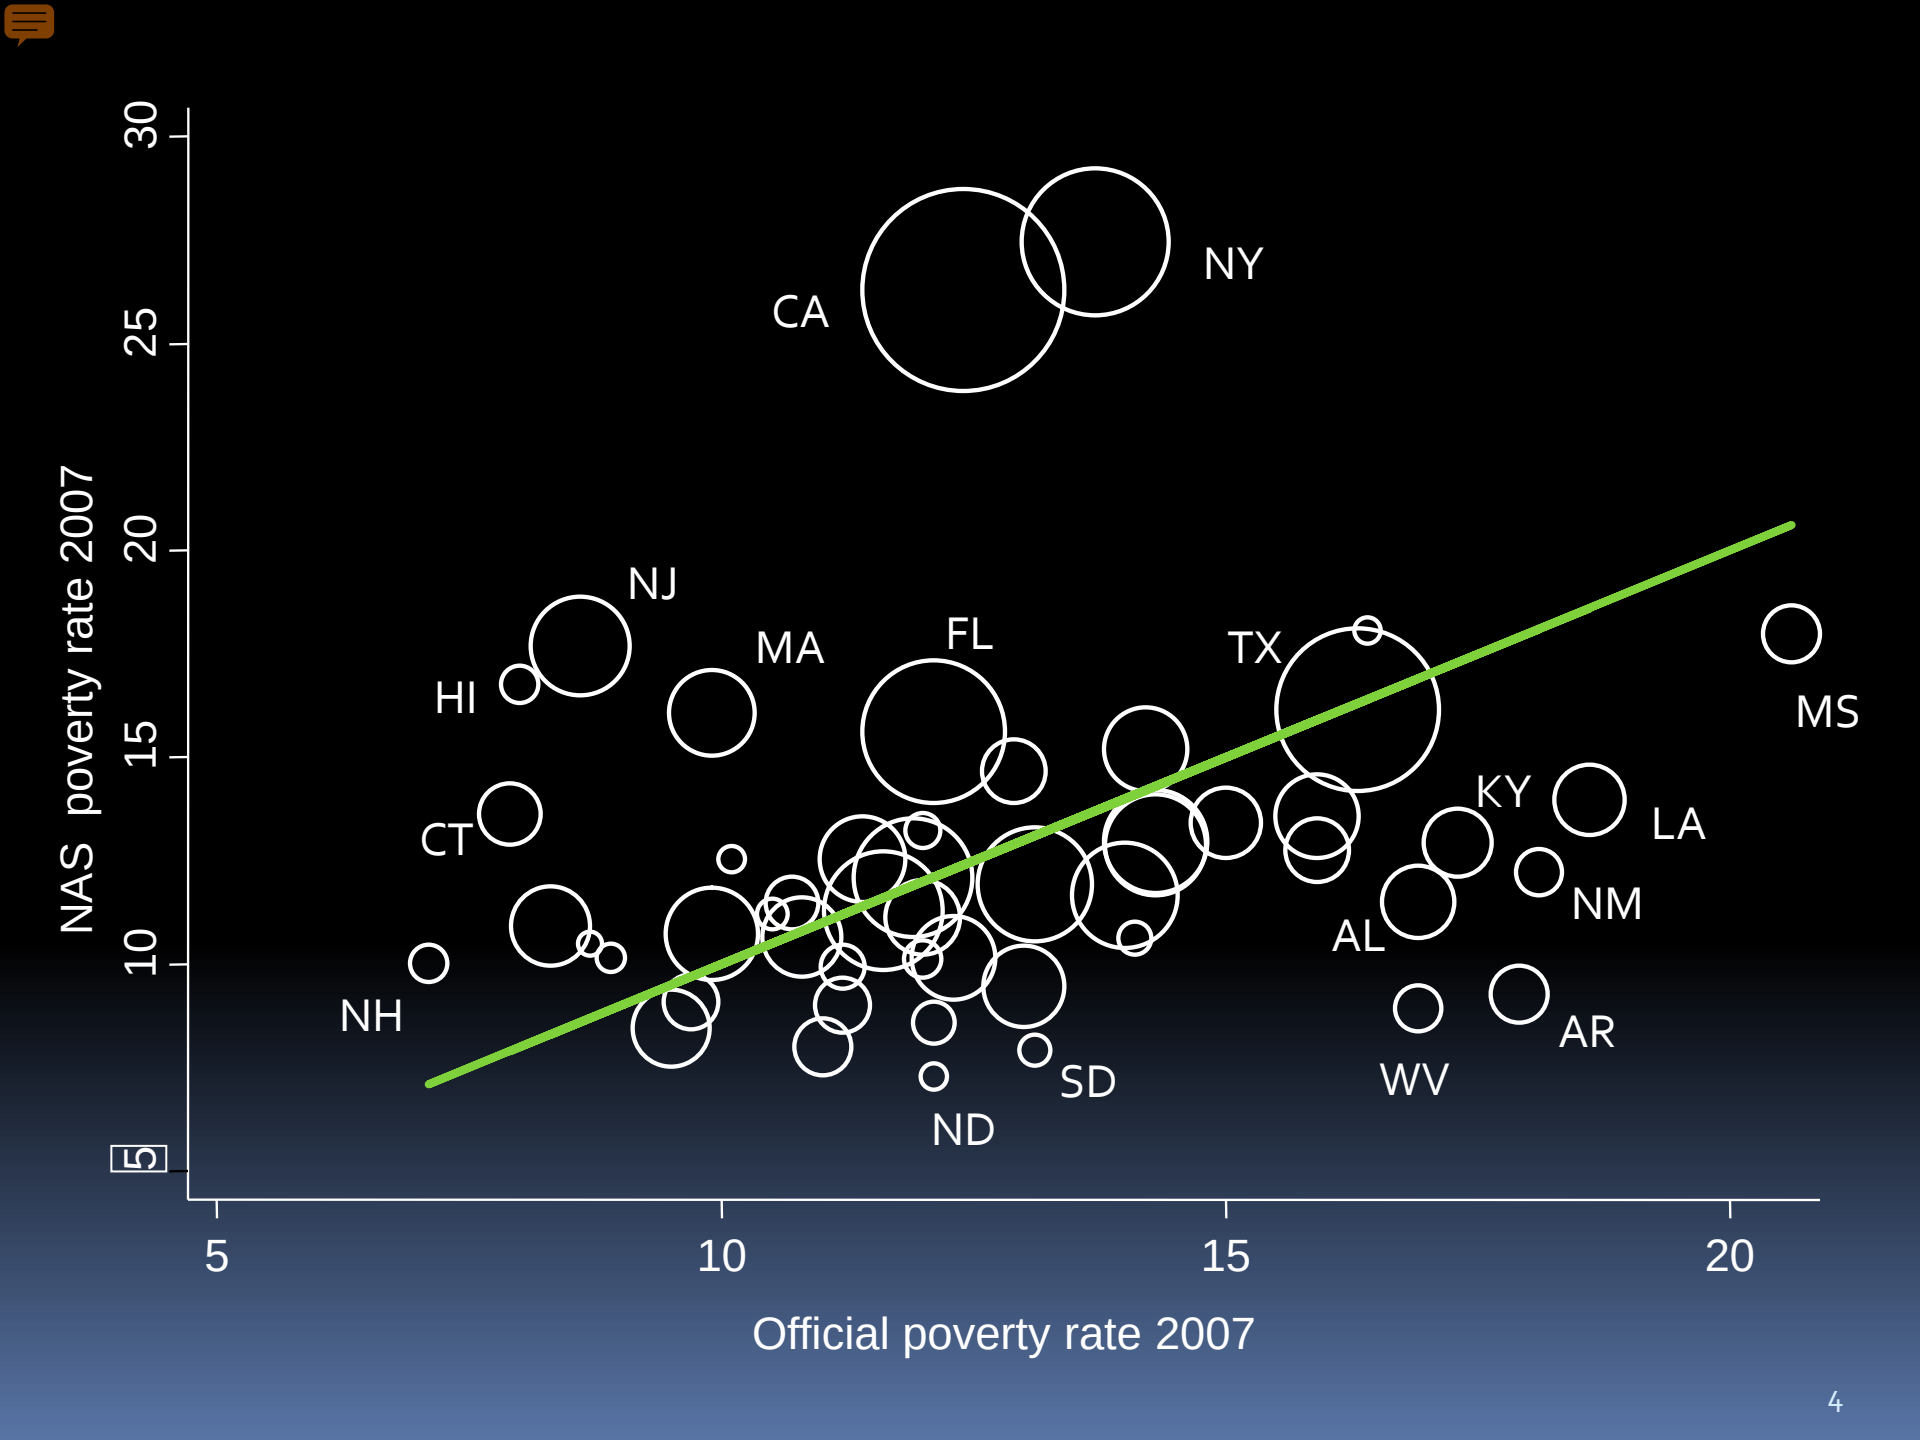
NAS poverty rate 2007
30
25
20
15
10
5
CA
NY
NJ
MA
FL
TX
MS
HI
CT
KY
LA
NM
AL
NH
AR
WV
SD
ND
5
10
15
20
Official poverty rate 2007

What I will talk about

- The ICP, global poverty, and things I learned
  - Which are relevant here
  - This material is often unfamiliar to economists who know a lot about price indexes
  - Yet it is important for spatial price indexes
    - Multilateral versus bilateral price indexes
- The current state of debate and measurement in the US
- Well-being, income and prices
  - In the world
  - In the US

Another qualification

- In the summer, I spend 5 weeks in south western Montana
- The rest of the year, I spend in NJ (most) or airplanes (too much)
- Hard not to compare MT and NJ and geographical prices and well-being
- Some contrasts

MT v NJ

- Aten et al say NJ RPP is 25.8 percent higher than MT RPP
- What would I earn in MT?
  - President Waded Cruzado of MSU gets $280,000
  - President Shirley Tilghman of Princeton gets $783,459
  - Ratio is 35.7 percent
  - Known in Bozeman, MT as "eating the scenery"
  - So I could expect a 2/3 pay cut
- I am comparing Montana in *August*
  - Montana is different from September to July
  - Seasonal effects may be important
- I doubt my well-being would be higher in MT, with a 55 percent real pay reduction
  - I might well live there if my tastes were different
  - People who live there *do* have different tastes
- I will return to these issues

Global poverty & the ICP

- There is a global poverty line ($-a-day), derived as the PPP converted average of domestic poverty lines of poorest countries
- Converted to domestic currencies at PPP for final consumption of households
  - This is a NIPA concept, rather than a CPI concept
  - Initial calculations used PPP for GDP
- Household surveys in each country are used to count numbers of people living in households whose per capita expenditure is below the line
  - Consumption, rather than an income standard
- Numbers added up to give a total of global poverty

Double role of PPPs

- PPPs from the ICP come into play *twice*, to calculate the line, and then to convert it to domestic currency
  - Indexes are very important, and small changes can have large effect
  - Density of global population is high near global poverty line
  - Especially in India, and to a lesser extent, in China
- Only second has parallel in US, but first might too? Why couldn't MT make its own poverty line?

International indices, PPPs

- Much less simple than they may seem
  - Price levels in each country sounds easy
- Fundamental tension: we need *both*
  - Goods that are strictly internationally comparable
  - Goods that are widely consumed in each country
    - Too loose: underestimate prices in poor countries
    - Too tight: overestimate prices in poor countries
- Large fractions of goods in domestic CPIs *excluded* from ICP
  - Domestically within the US, these issues are most important for housing
    - What is it?
    - How do we control for quality differences?
    - What sort of unit do we specify?
    - But not totally absent from other goods too

Weights for PPPs

- The ICP collects prices on thousands of goods and services
  - But there are expenditure weights for 102 "basic heads" of consumption, from NIPA
  - Used as weights at an "upper stage"
- In the US, we do not have this detail for state or city accounts
  - See Aten et al for the best that can be done
  - In this respect, US is less strong than ICP

Deeper issues with weights

- PPPs are *multilateral* indexes, not bilateral indexes
  - Lots of flavors: a simple one is to calculate Fisher (or other symmetric) price index for each *pair* of countries (or states)
  - Enforce transitivity by some sort of adjustment (Gini suggested least squares)
  - So that price in NJ relative to MT are the same if calculated directly, or along I80, via WY, NE, IA, IL, IN, OH, and PA
    - Or any other chain
- Transitivity and symmetry come at a price
  - Can't price a fixed bundle of goods: no Laspeyres
  - Consumption patterns in FL affect price index for NJ relative to MT; e.g. if more older people move to FL

Air travel in Cameroon

- In some calculations, Cameroon was more expensive than Hong Kong in 2005
- Air travel is *very* expensive in Cameroon
  - Tiny budget share of air travel in Cameroon
- Need a pairwise index with UK, for example
  - Budget share in UK is 2 x share in Cameroon
  - Price is 11.6 times average ratio
- Symmetric index effectively averages the weights
  - So weight applied to 11.6 is half of Cameroon and half of UK
  - Largest BH contributor to PPP for Cameroon v UK
  - British consumption of air travel makes more poor people in Cameroon
- This could occur in the US for housing in a comparison of SF or NY with mid-West rents
  - About three times, not 12 times
  - But issue remains of whose weights to use

COLI v Price Index

- If everyone had the same tastes, many of these issues would evaporate
  - Though still problem of using Cameroon or British utility level
  - This would vanish too if tastes were identical and homothetic
  - Some are happy with this assumption
- Perhaps problems less in the US
  - To some extent
  - But not if people locate based on their tastes
  - Or if geography affects tastes
  - MT versus NJ again

Weights for the poor?

- Criticism of global poverty calculations is that weights are NIPA weights
  - Plutocratic price indexes
  - Poor have very different bundles
- Olivier Dupriez and I investigated this
  - Replaced NIPA by household surveys
  - Calculated weights for those near the global poverty line
  - Makes remarkably little difference
- Reasons are explored in our paper (AEJA, April 2011)
  - True that poor have very different consumption patterns
  - But what matters is whether reweighting has differential effects in different countries
  - As well as structure of relative prices
  - Hard to know without doing it, but unlikely to be first-order effect in the US

Politics of the ICP

- Countries have an incentive to overstate their prices if it induces transfers or for other reasons
  - This is easy to do if countries have control over price collection
  - Not true in US states now, but could become a political issue
  - See it already in discussions of what kind of housing to price
- No one much cares about global poverty
  - No transfers depend on it, or the ICP
  - For some, more is better for their funding
  - Undermines the statistical credibility of the calculations
  - When WB mistakenly increased the numbers of global poor by 50 percent, and blamed it on the ICP revision, very little reaction

Europe and India

- Better parallel for US is Eurostat
  - Calculates PPPs for Euroland every year
  - Those numbers are used for setting transfers
  - They *do* contradict national CPIs
  - They *do not* use a common market basket
  - These numbers have high political salience and are carefully protected
  - Calculated by high quality, well-organized statistical service
  - With protections against political interference
- India indexes its poverty lines by state and urban/rural within each

The big difference!

- People in India or Tanzania are not free to move to Japan or the US
- People in MT can move to NJ, or vice versa
- One could *imagine* an equilibrium in which each person is located where their well-being is maximized
  - People in MT like to "eat the scenery" and will pay for it
  - People in NJ like to visit NY city, go to baseball games, and eat in nice restaurants, buy fresh fish
  - Rosen/Roback framework
  - Unclear that we want to index benefits to local prices that reflect amenities
- This model is unrealistic in various ways, but it surely the correct starting point for investigation
  - I think now that the NAS panel got this wrong, or at least did not do the work and make the arguments that need to be made

Special case

- Country that exports oil/coffee and imports all consumer goods
- Prices are higher the further from the port
- Incomes adjust to those prices because of mobility
- Deflating by local prices is the right thing to do
- But to make this argument we need to know *why* prices are higher in some places than others
- Still might need to adjust for local amenities

COLI v RPP

- Old question of what counts in utility and is supposed to be compensated for by COLI in a price-like way
  - What does the cost of living include, beyond prices?
  - Why are we making these calculations?
  - For which benefits?
- Cold in MT in winter
  - Even if prices were the same, COLI would be higher
  - This one seems OK: close to higher prices
- The value of amenities may change if income falls or be different for different people
  - Or for the poor and non-poor

COLI v RPP (2)

- Life expectancy is higher in UT than in NV
  - Should this be embodied into a COLI?
  - I doubt it: not like lower prices
  - Doesn't really help a poor person to know that life if long
- Rents may be higher
  - But choice is not only on rents but on commuting time
  - Very important joint decision
  - NAS panel looked at this, but was nervous because much commuting expenditure is a luxury choice
- I know of no principled way of settling these
  - Which is why I argued for *price* indexes not *cost of living* indexes in NAS price report, because I at least know what a price index is

Maybe we can just ask?

- NAS panel was intrigued by Gallup's question about income needed just to get by
- Mean/median very close to Orshanky's original line
  - Which may have had more to do with this than with the food bundle
- Changed in line with median income since
- But we had no one with expertise in such concepts
  - We were nervous about basing a PL on polling
  - US statistical service more opposed to subjective data than is currently the case
  - Same among economists
- So a preliminary, very rough, exploration follows

Gallup World and Daily Polls

- World Poll, since 2006, polls 1,000 people in each country in each year
- Daily Poll, since Jan 2008, polls 1,000 people in the US each day
  - More than 1,000,000 observations now
  - Lots of geographical detail possible
- Range of well-being questions
  - Ladder (worst and best life, 0 to 10)
  - Standard of living (Satisfactory? Yes/no)
- Also single question on income
- Lots of demographics

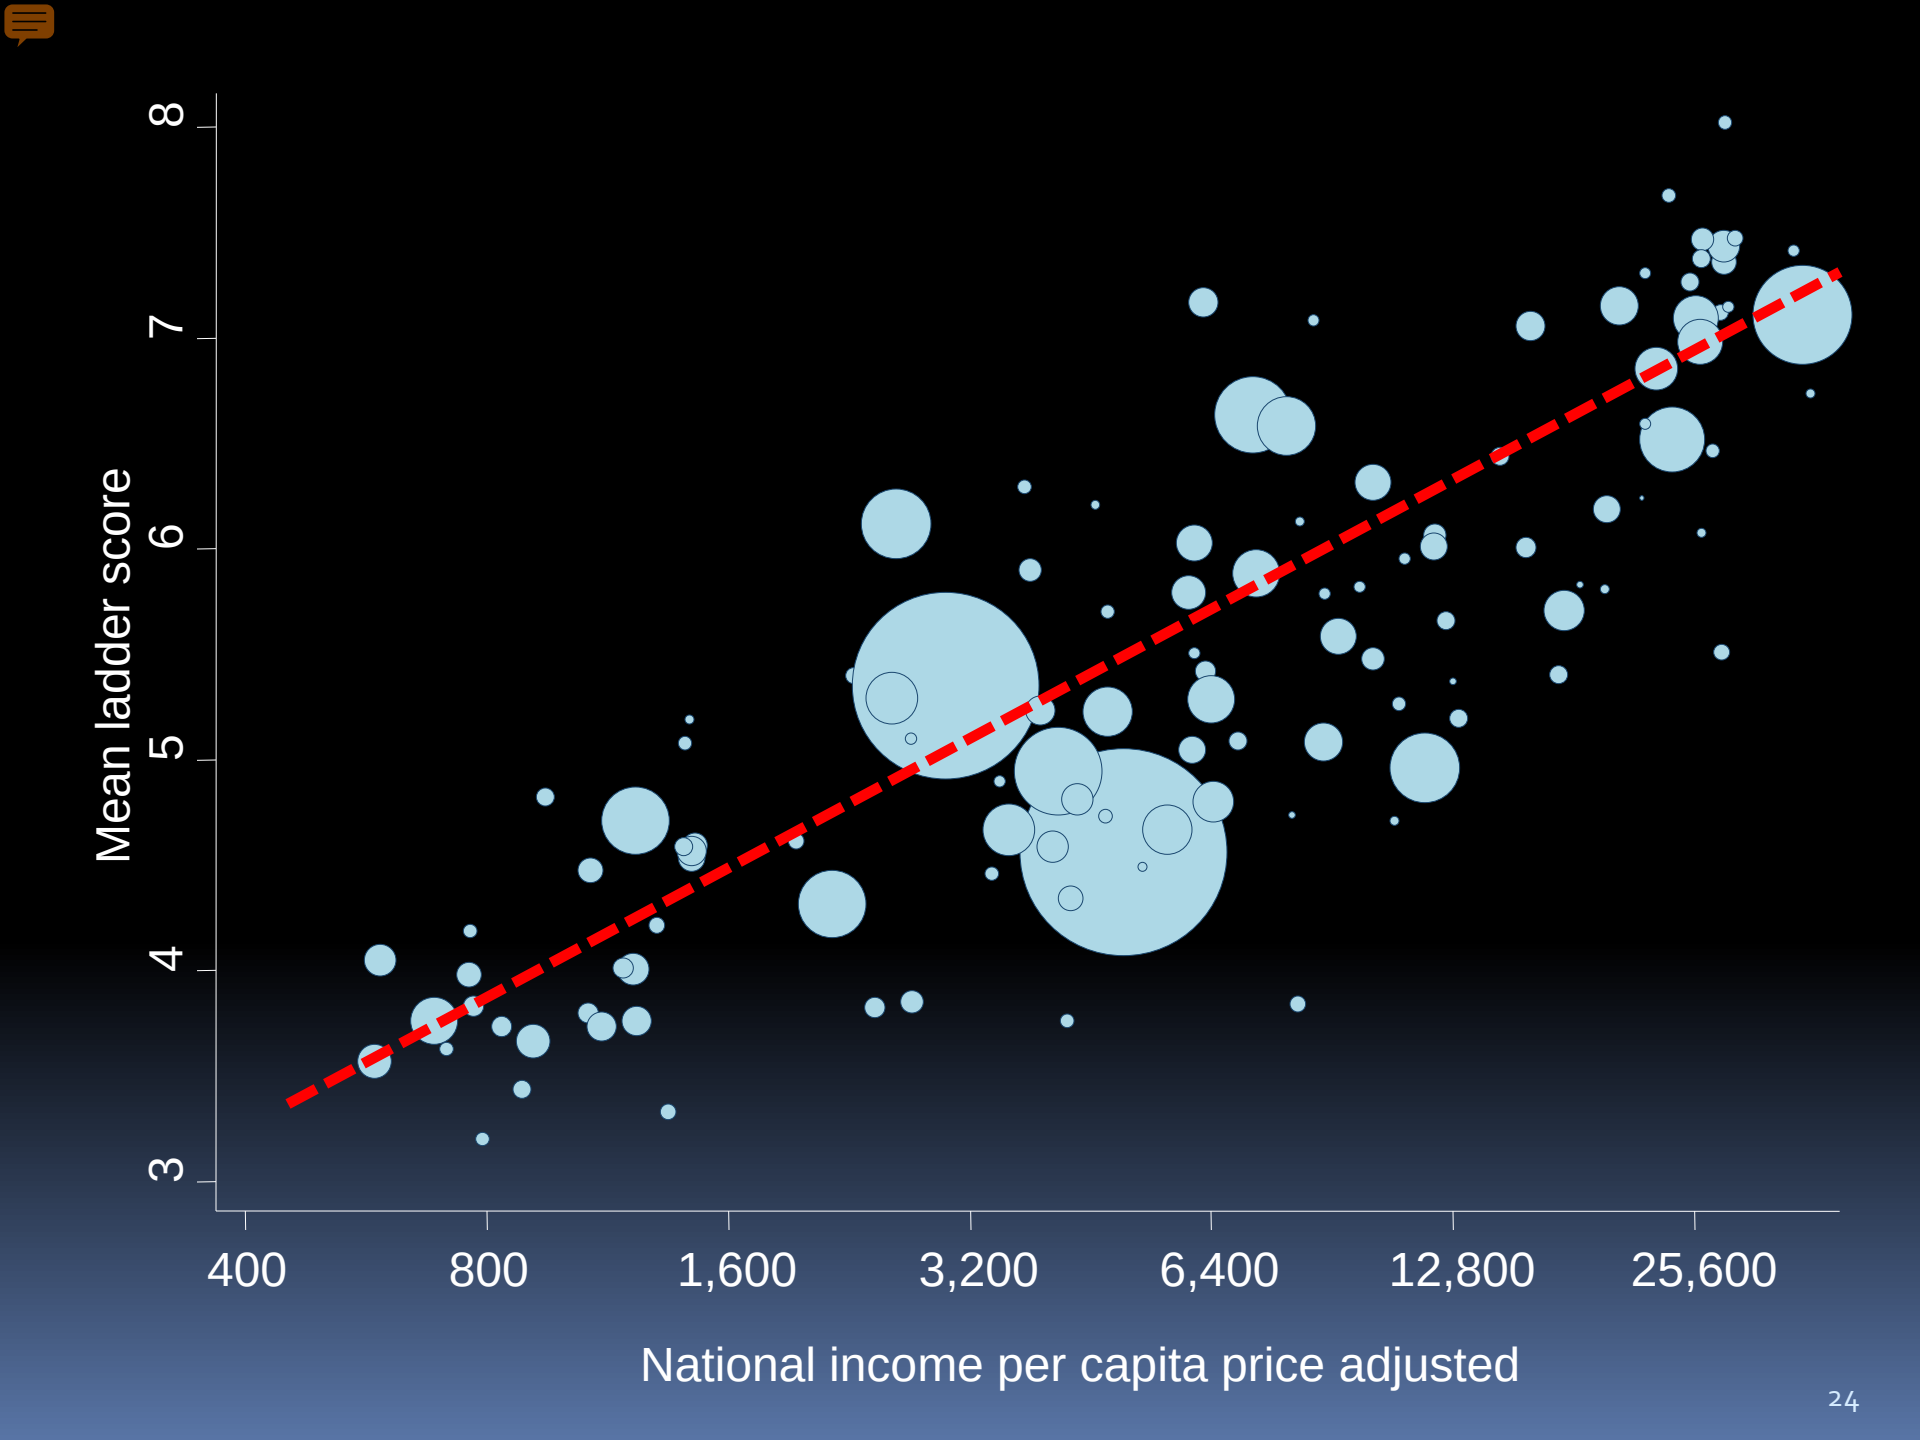
Mean ladder score
3
4
5
6
7
8
400
800
1,600
3,200
6,400
12,800
25,600
National income per capita price adjusted

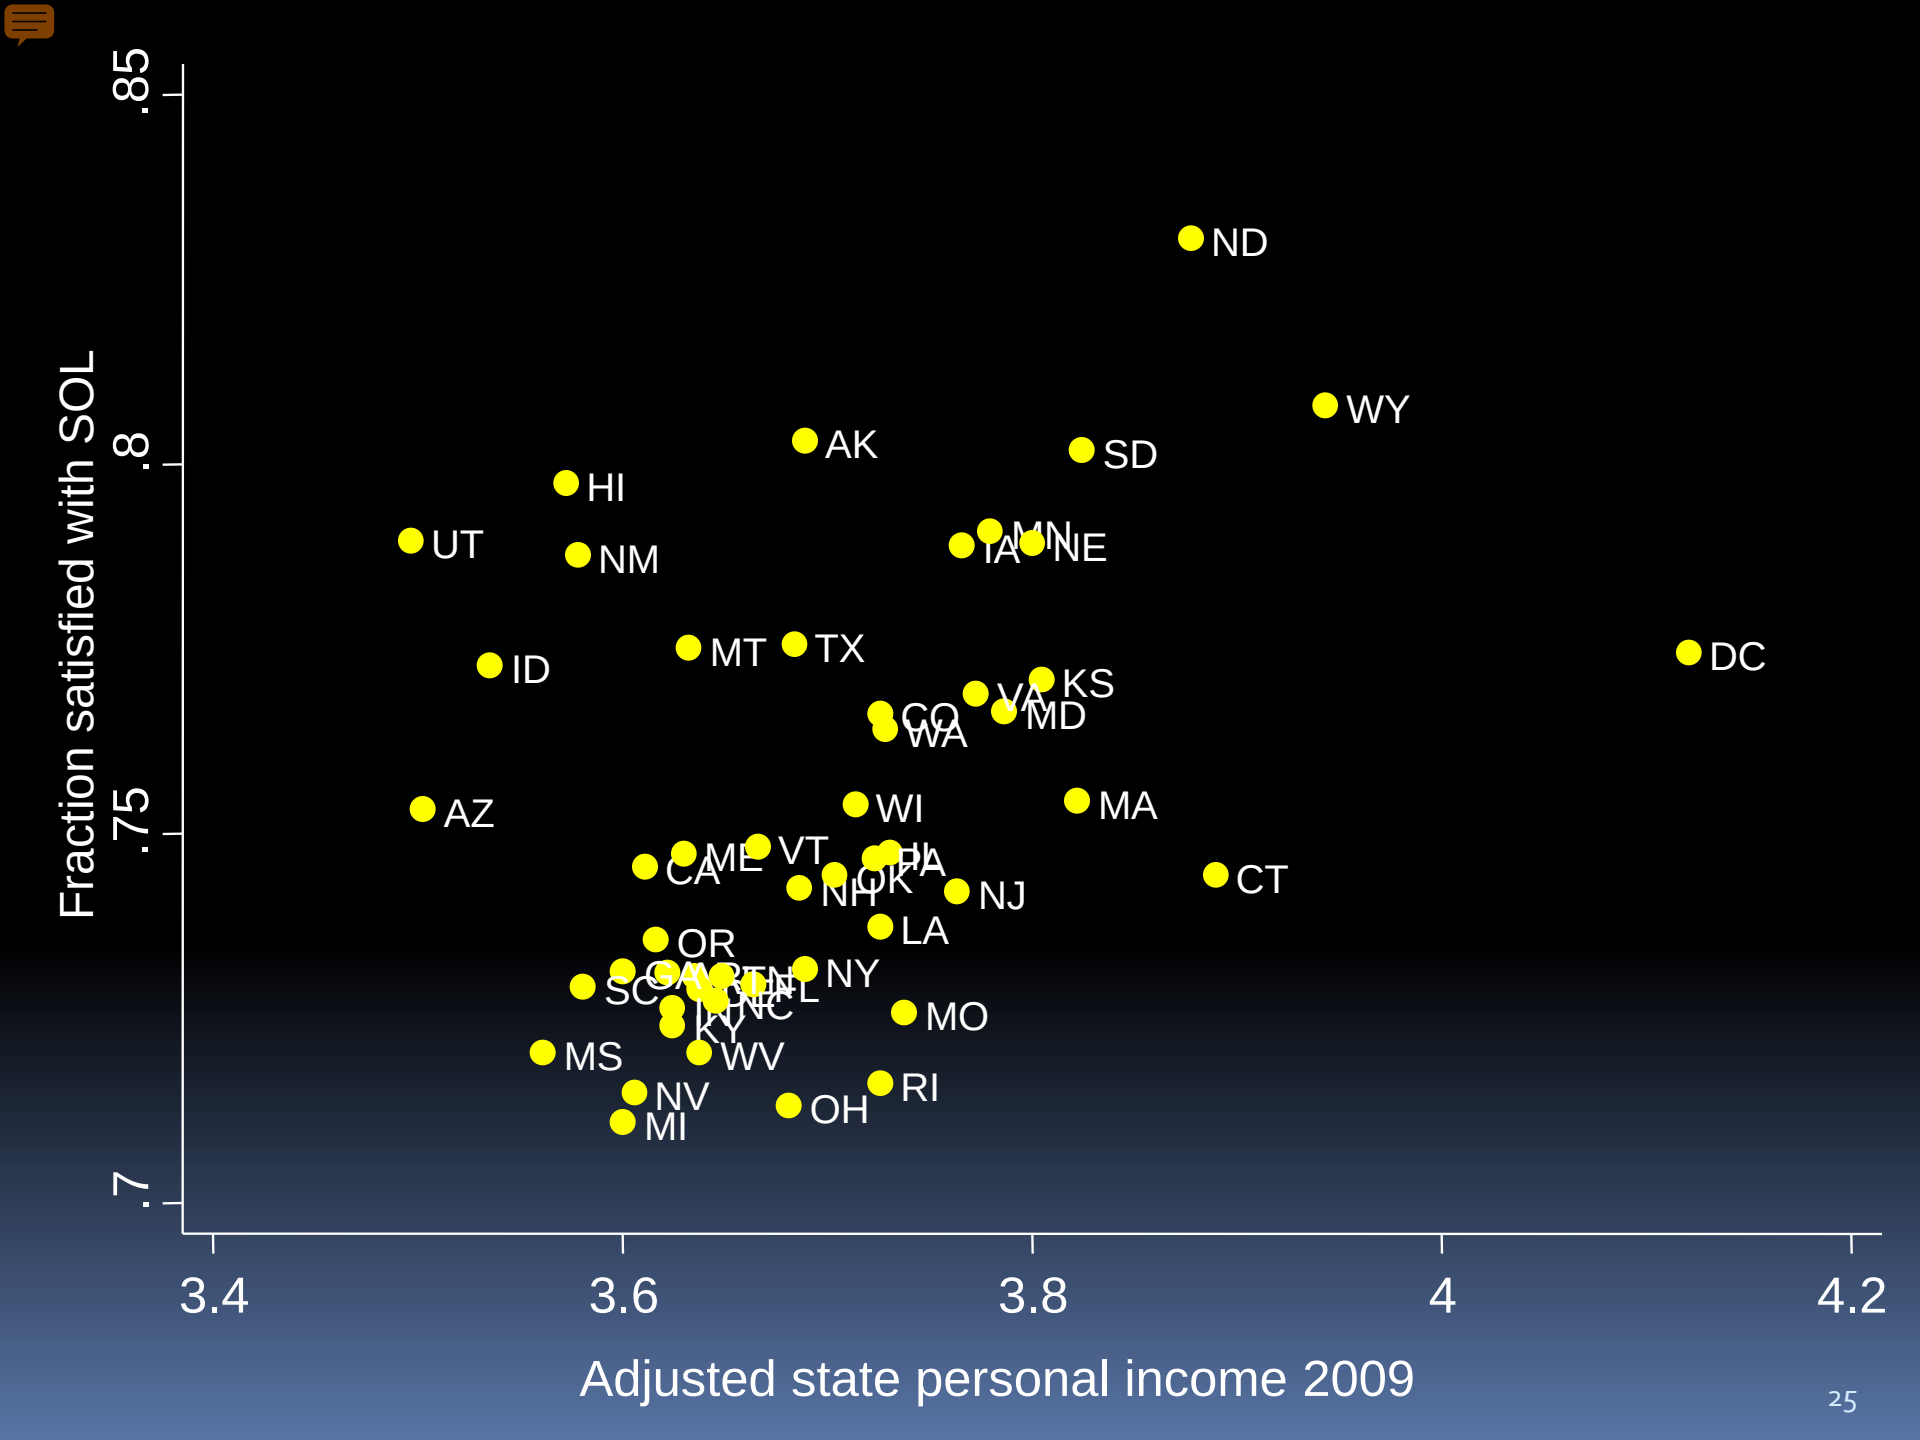
Fraction satisfied with SOL
.85
.8
.75
.7
3.4
3.6
3.8
4
4.2
Adjusted state personal income 2009
ND
WY
AK
SD
HI
UT
NM
IA
NE
MT
TX
DC
ID
KS
VA
MD
CO
WA
AZ
WI
MA
VT
CA
NH
NJ
CT
LA
OR
NY
MO
MS
WV
RI
NV
OH
MI

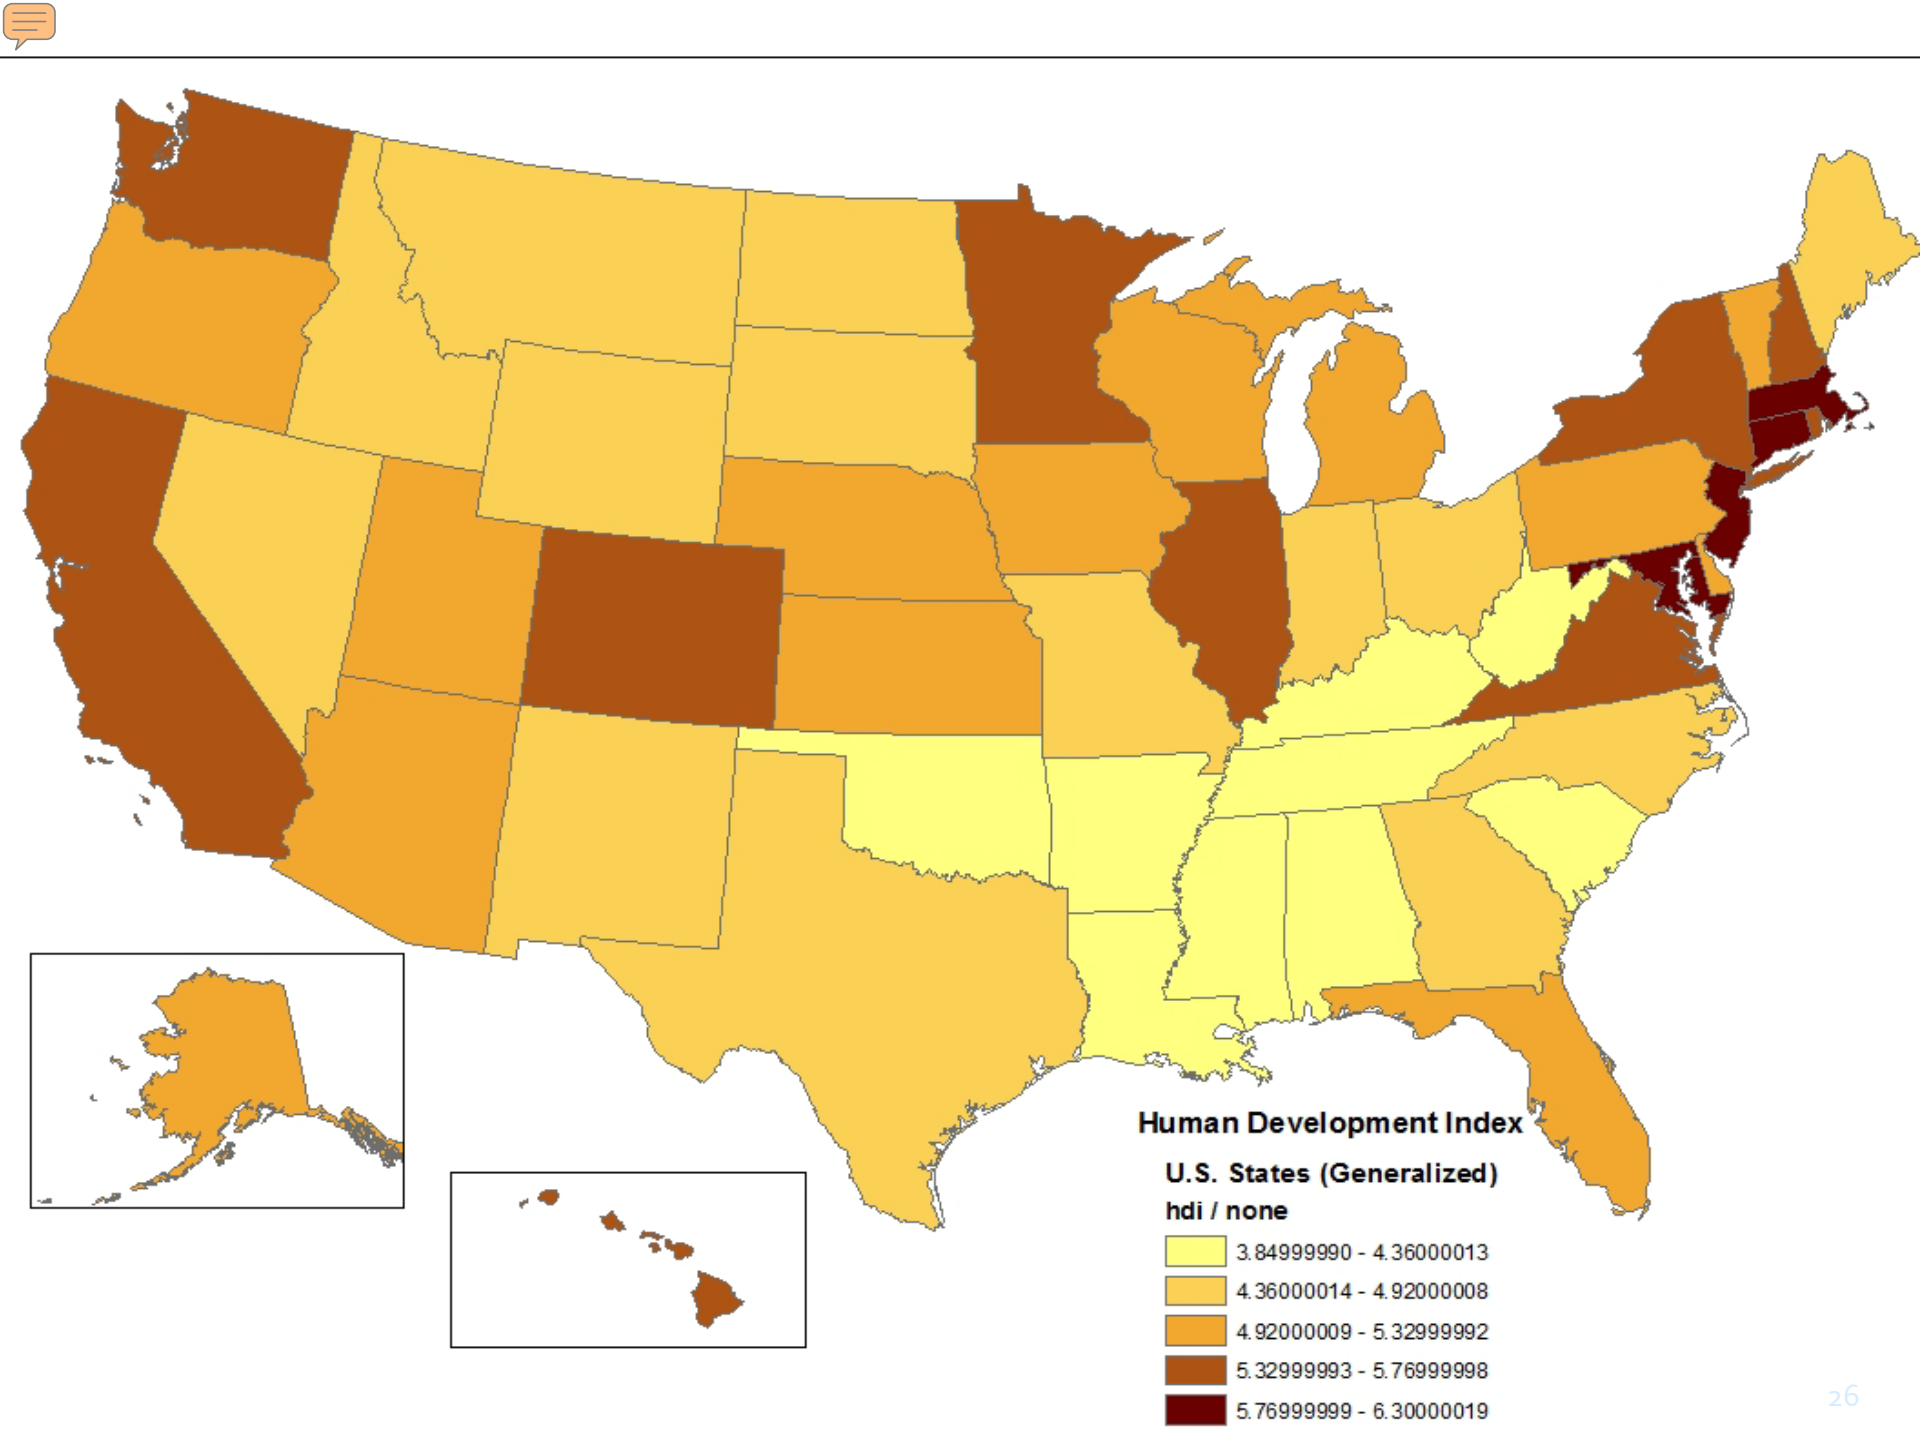
Human Development Index
U.S. States (Generalized)
hdi / none
3.84999990 - 4.36000013
4.36000014 - 4.92000008
4.92000009 - 5.32999992
5.32999993 - 5.76999998
5.76999999 - 6.30000019

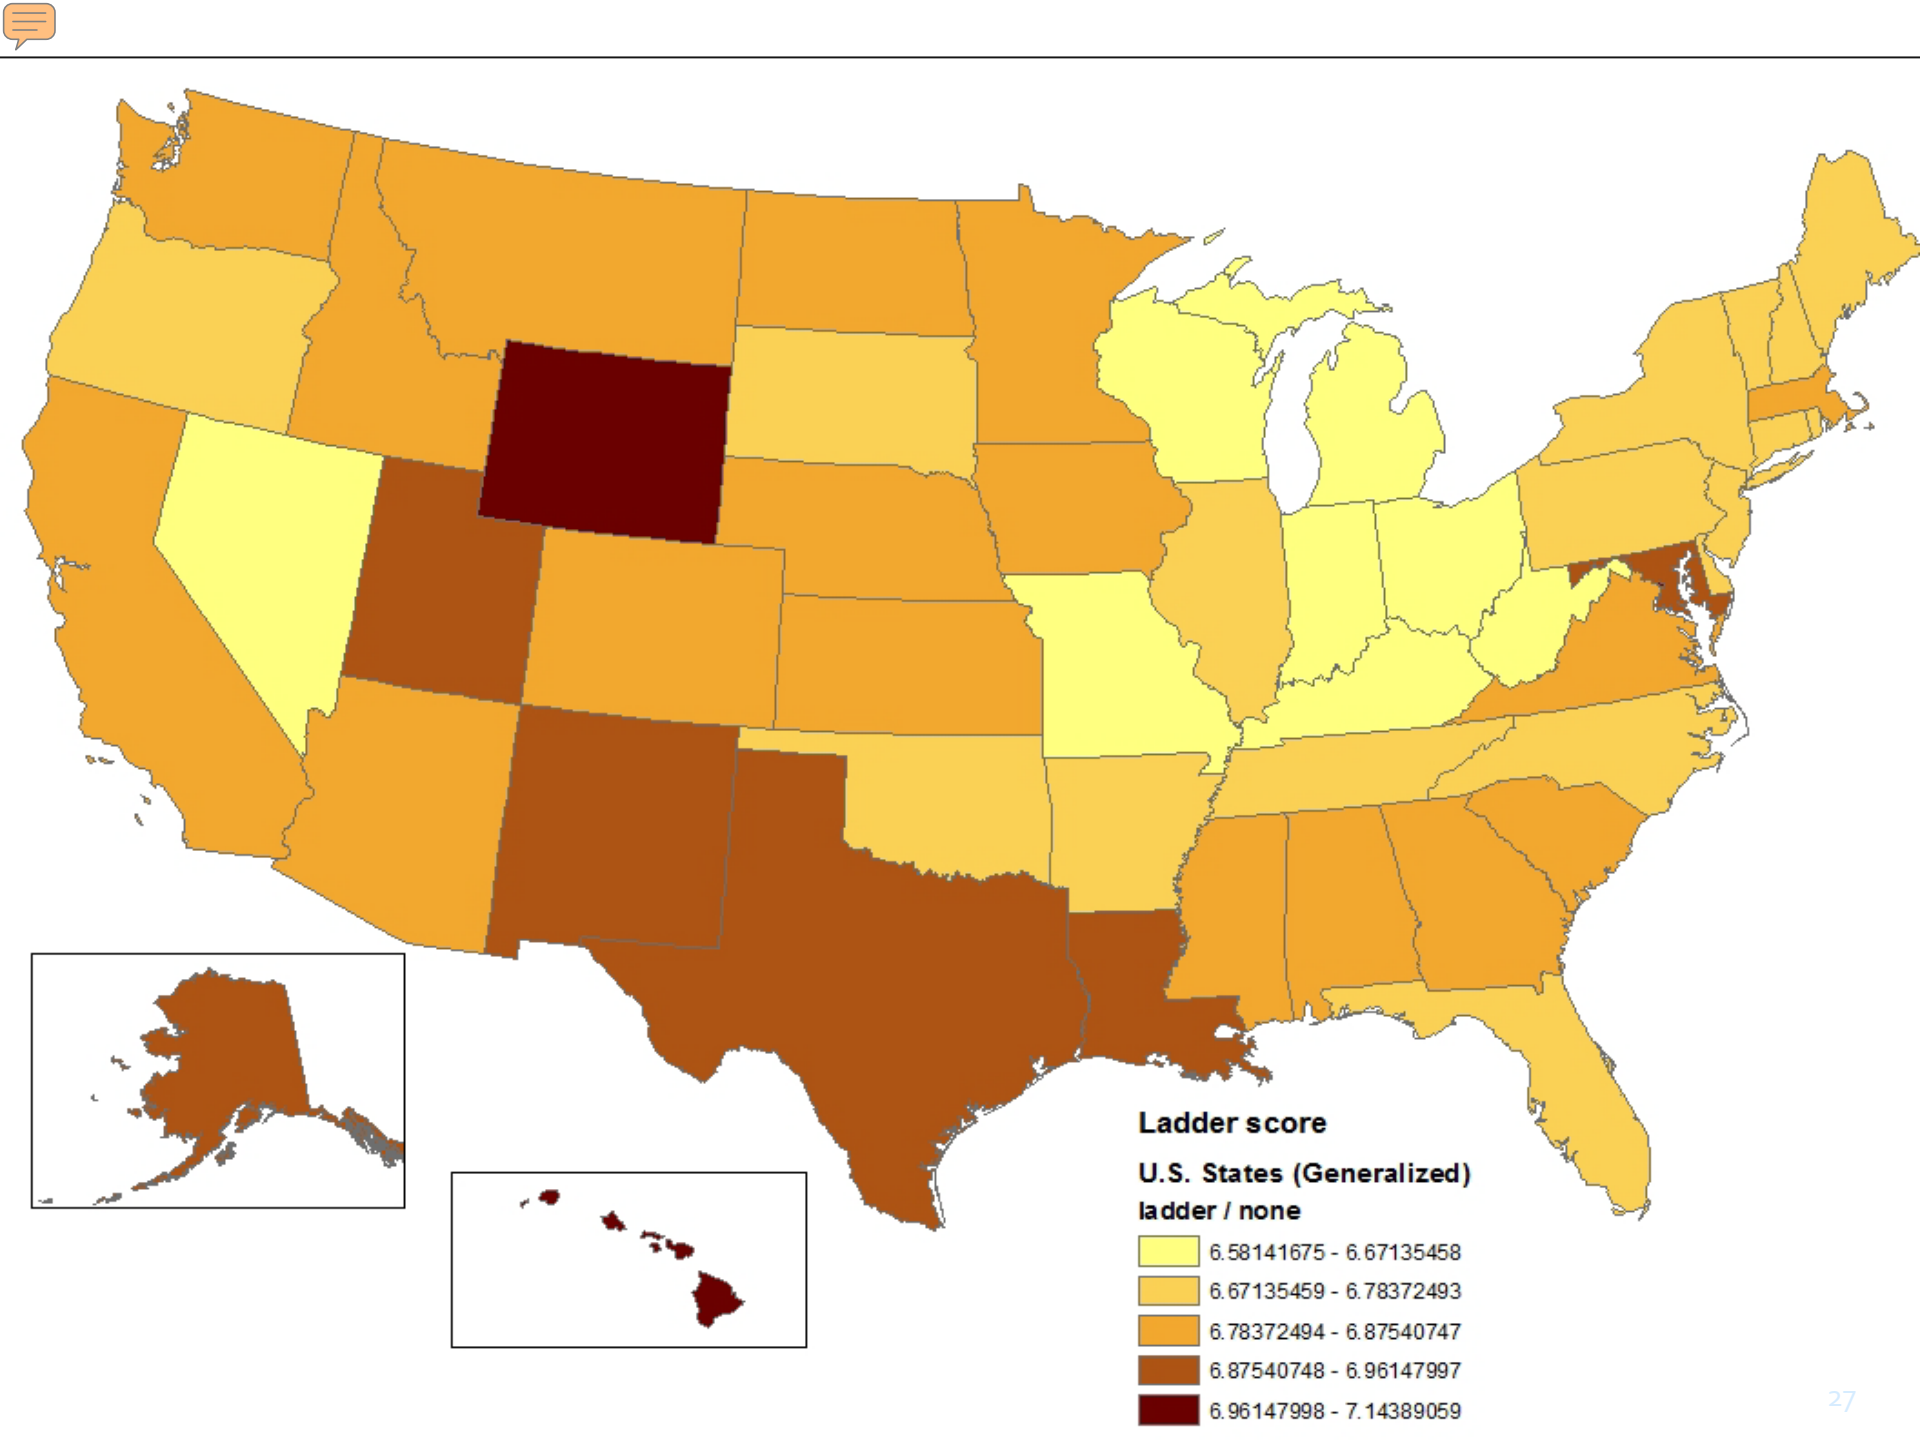
Ladder score
U.S. States (Generalized)
ladder / none
6.58141675 - 6.67135458
6.67135459 - 6.78372493
6.78372494 - 6.87540747
6.87540748 - 6.96147997
6.96147998 - 7.14389059

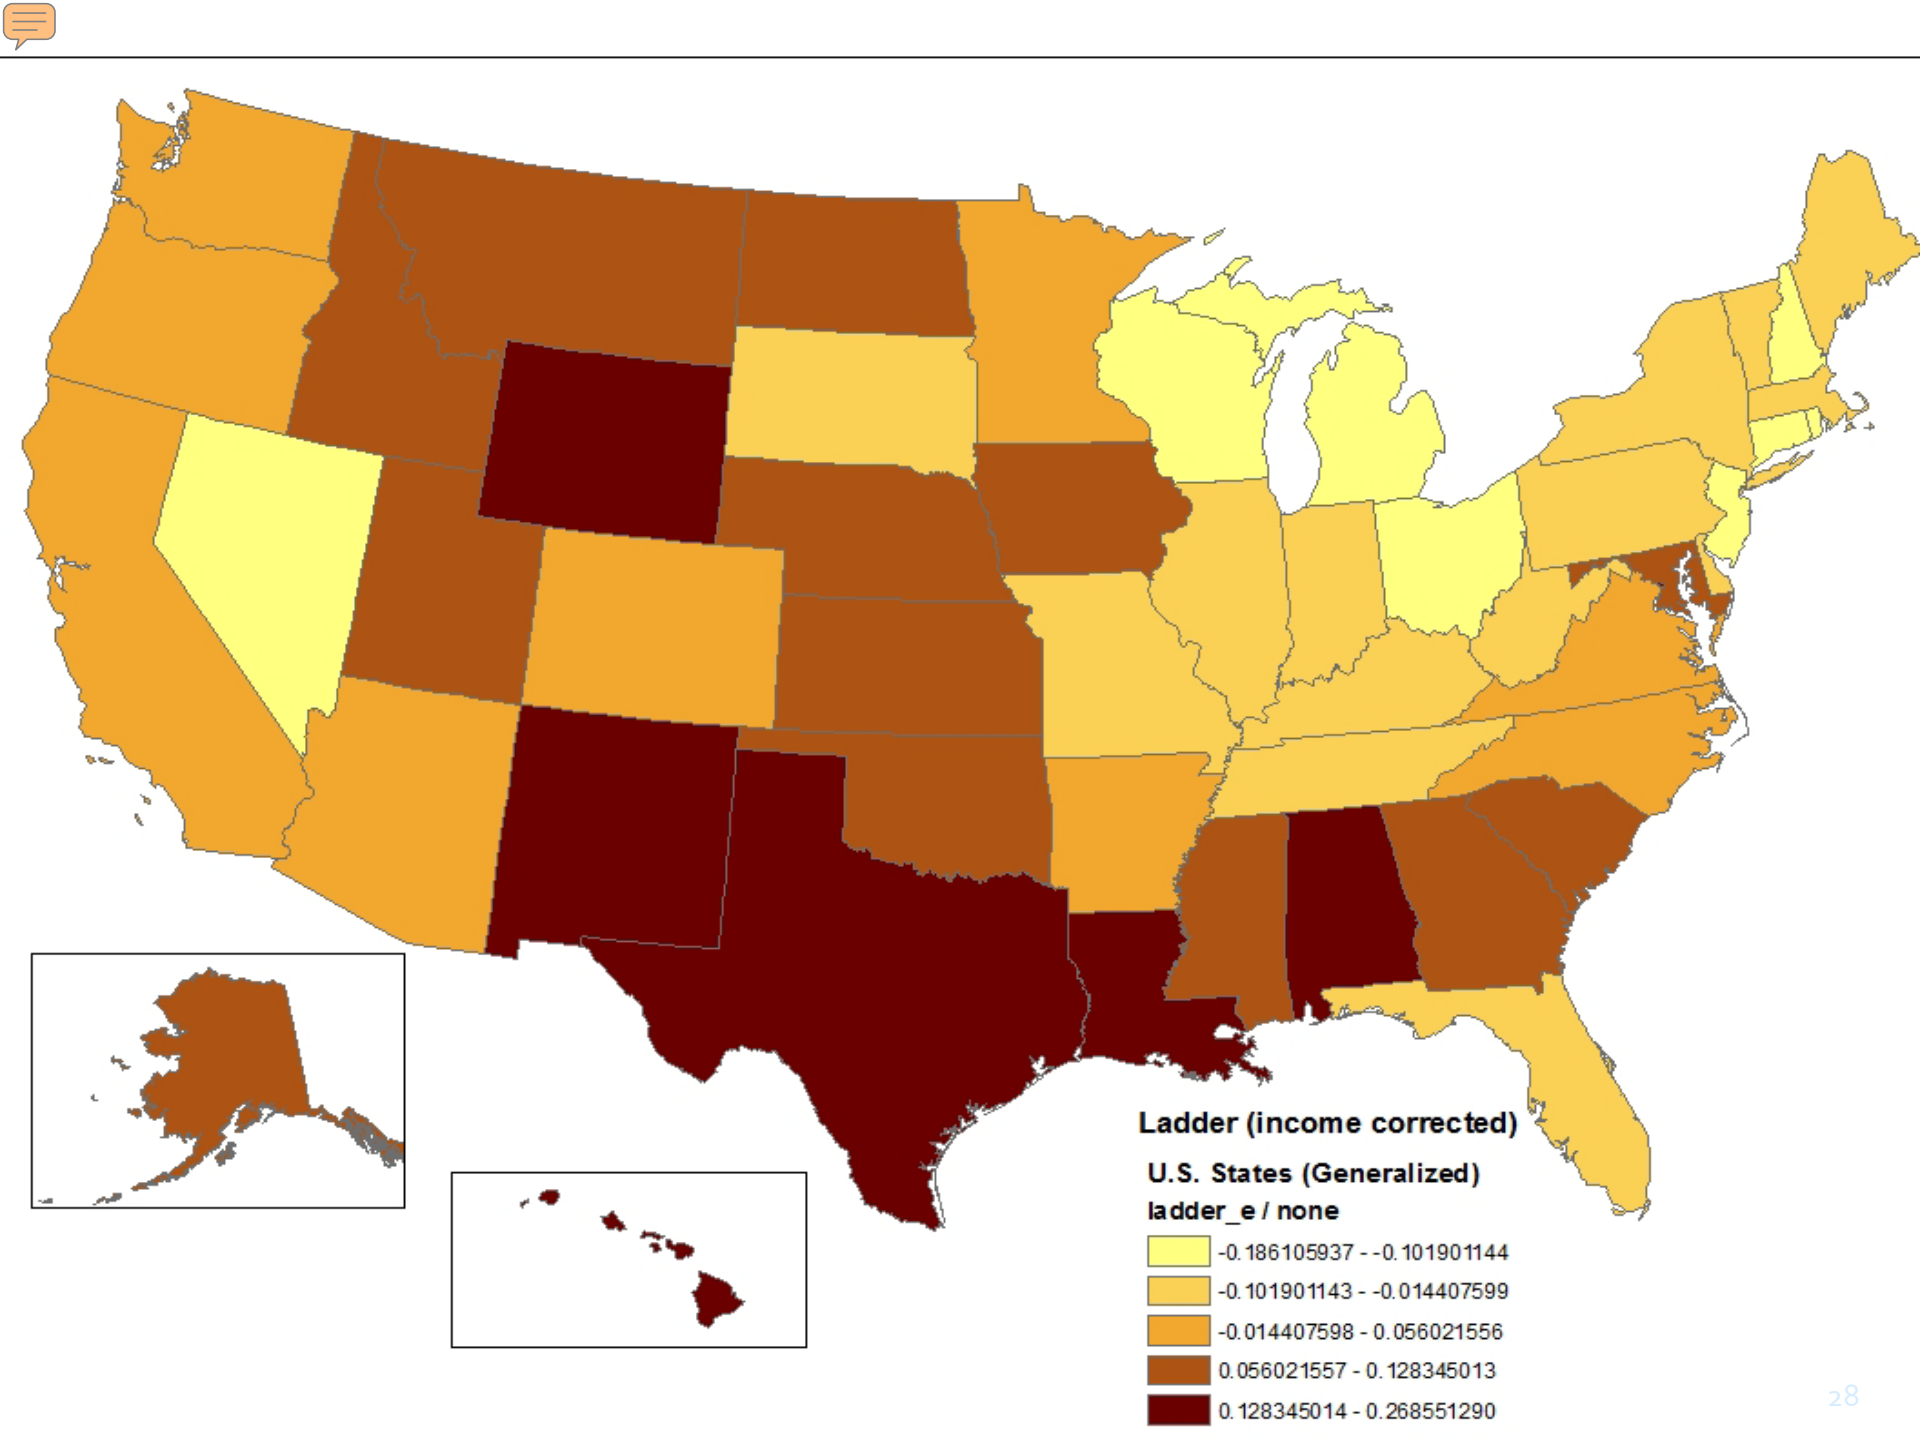
Ladder (income corrected)
U.S. States (Generalized)
ladder_e / none
-0.186105937 - -0.101901144
-0.101901143 - -0.014407599
-0.014407598 - 0.056021556
0.056021557 - 0.128345013
0.128345014 - 0.268551290

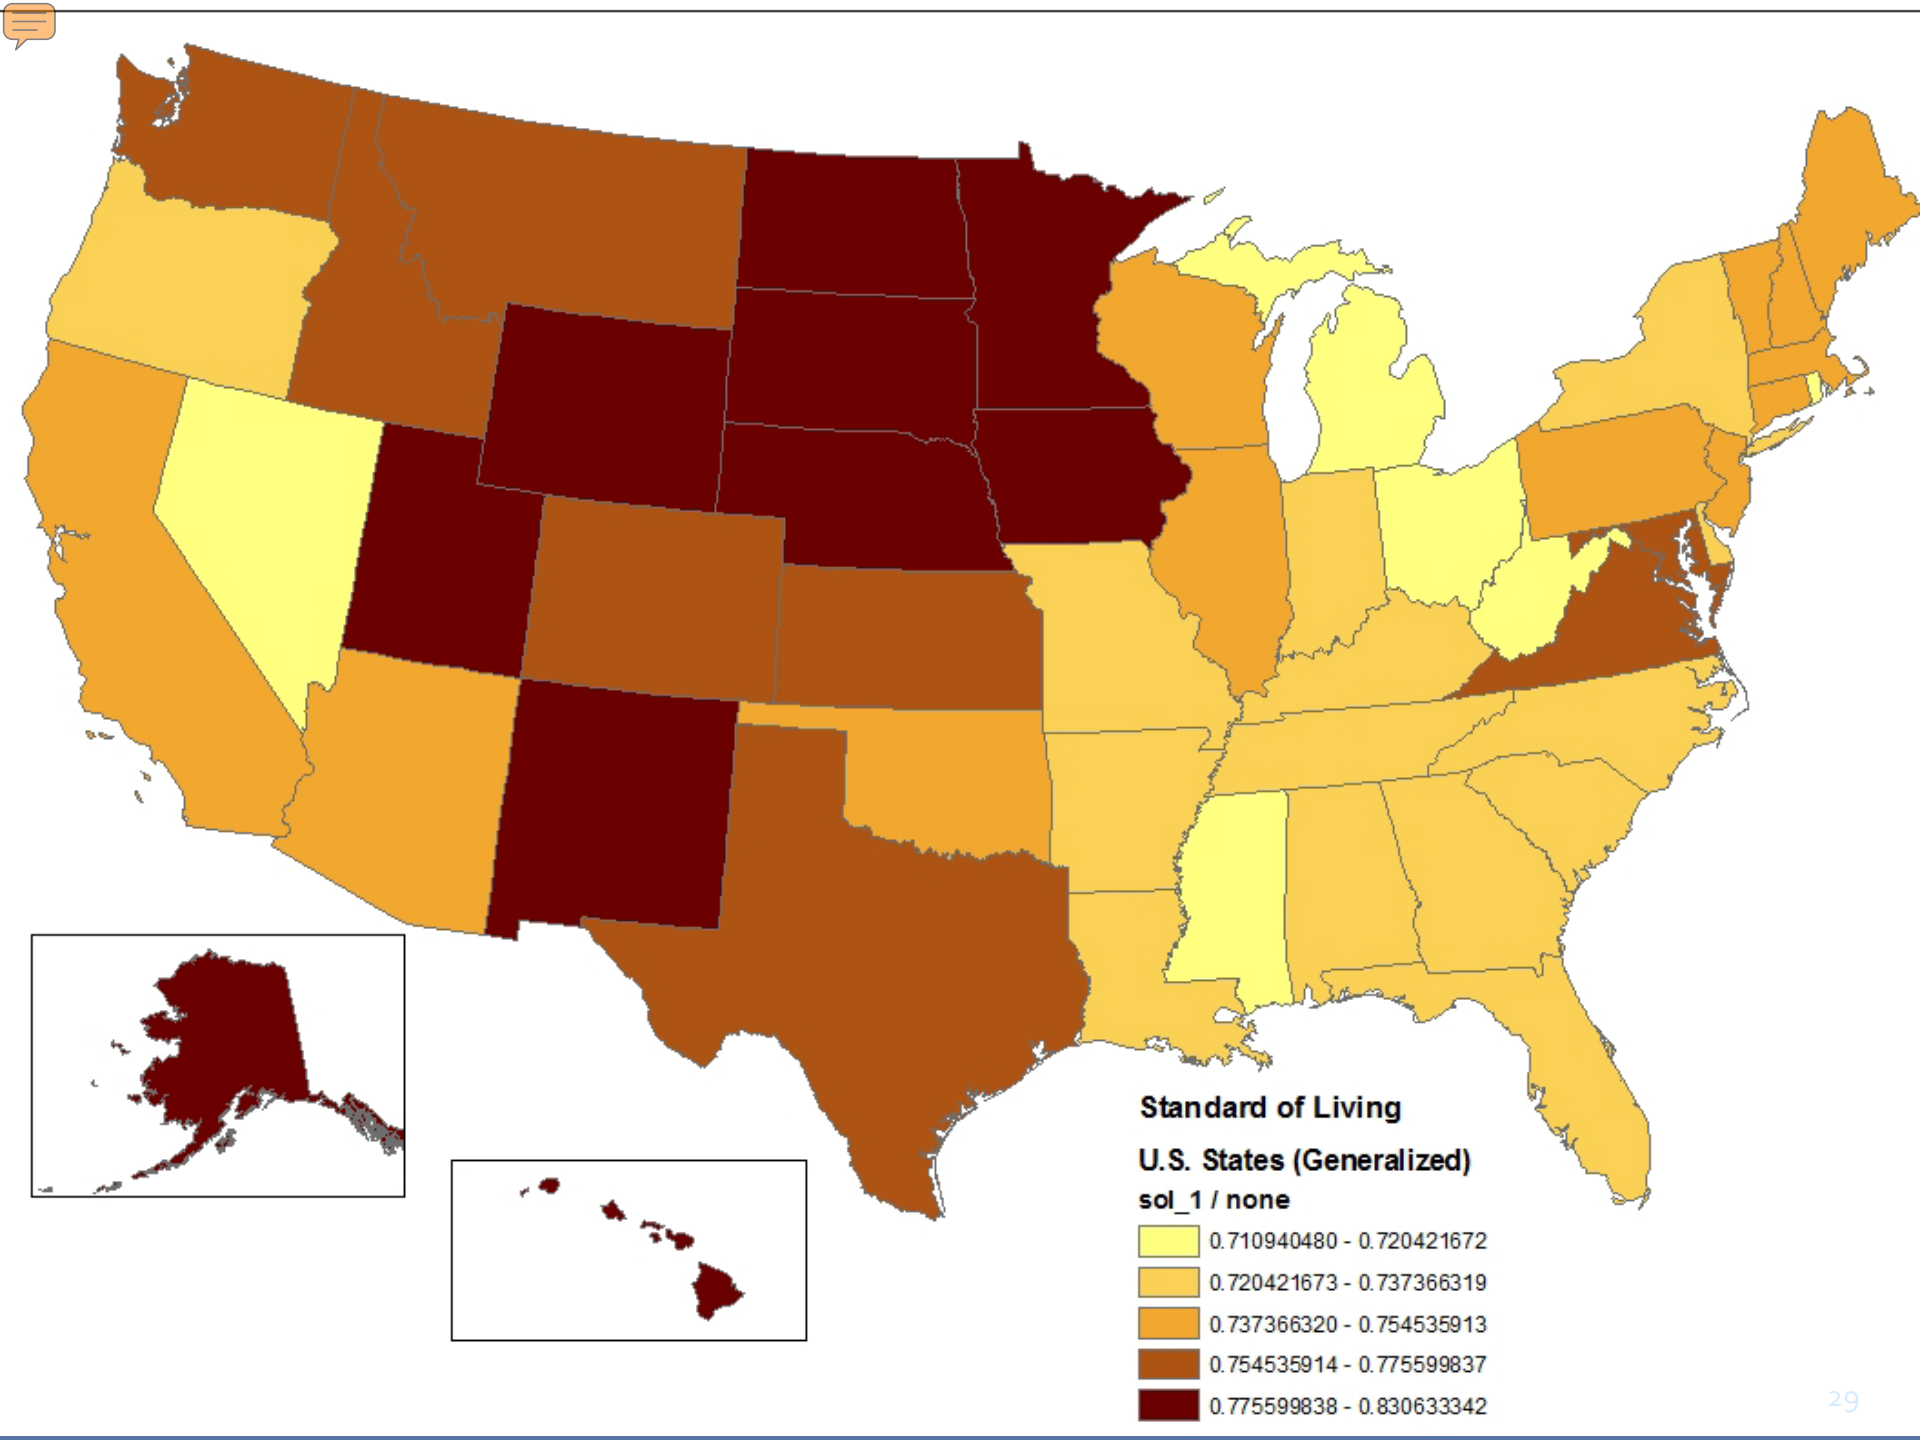
Standard of Living
U.S. States (Generalized)
sol_1 / none
0.710940480 - 0.720421672
0.720421673 - 0.737366319
0.737366320 - 0.754535913
0.754535914 - 0.775599837
0.775599838 - 0.830633342

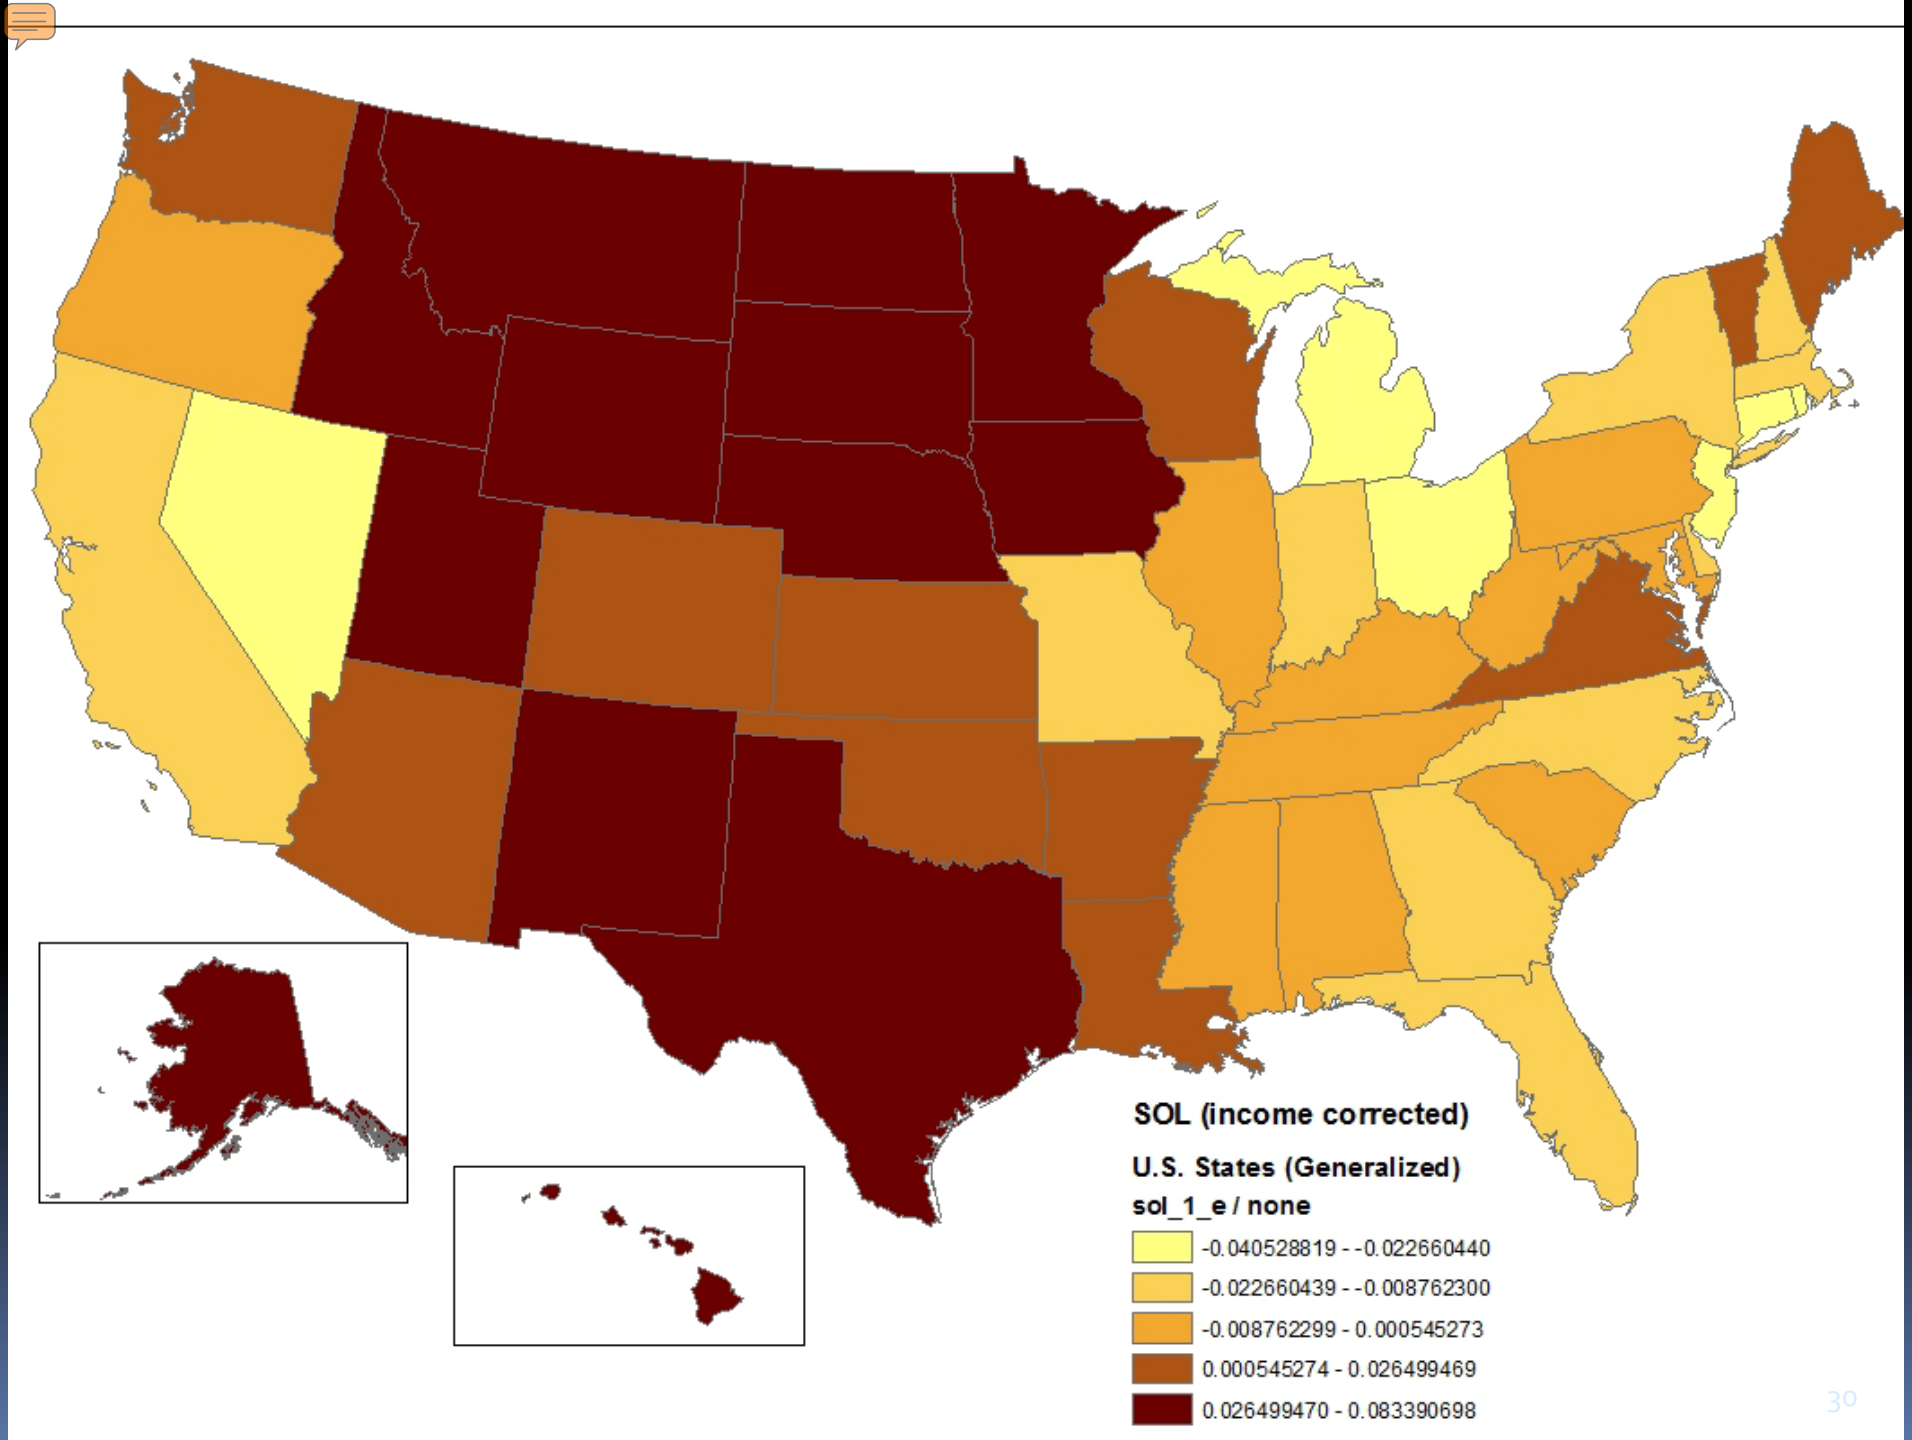
SOL (income corrected)
U.S. States (Generalized)
sol_1_e / none
-0.040528819 - -0.022660440
-0.022660439 - -0.008762300
-0.008762299 - 0.000545273
0.000545274 - 0.026499469
0.026499470 - 0.083390698

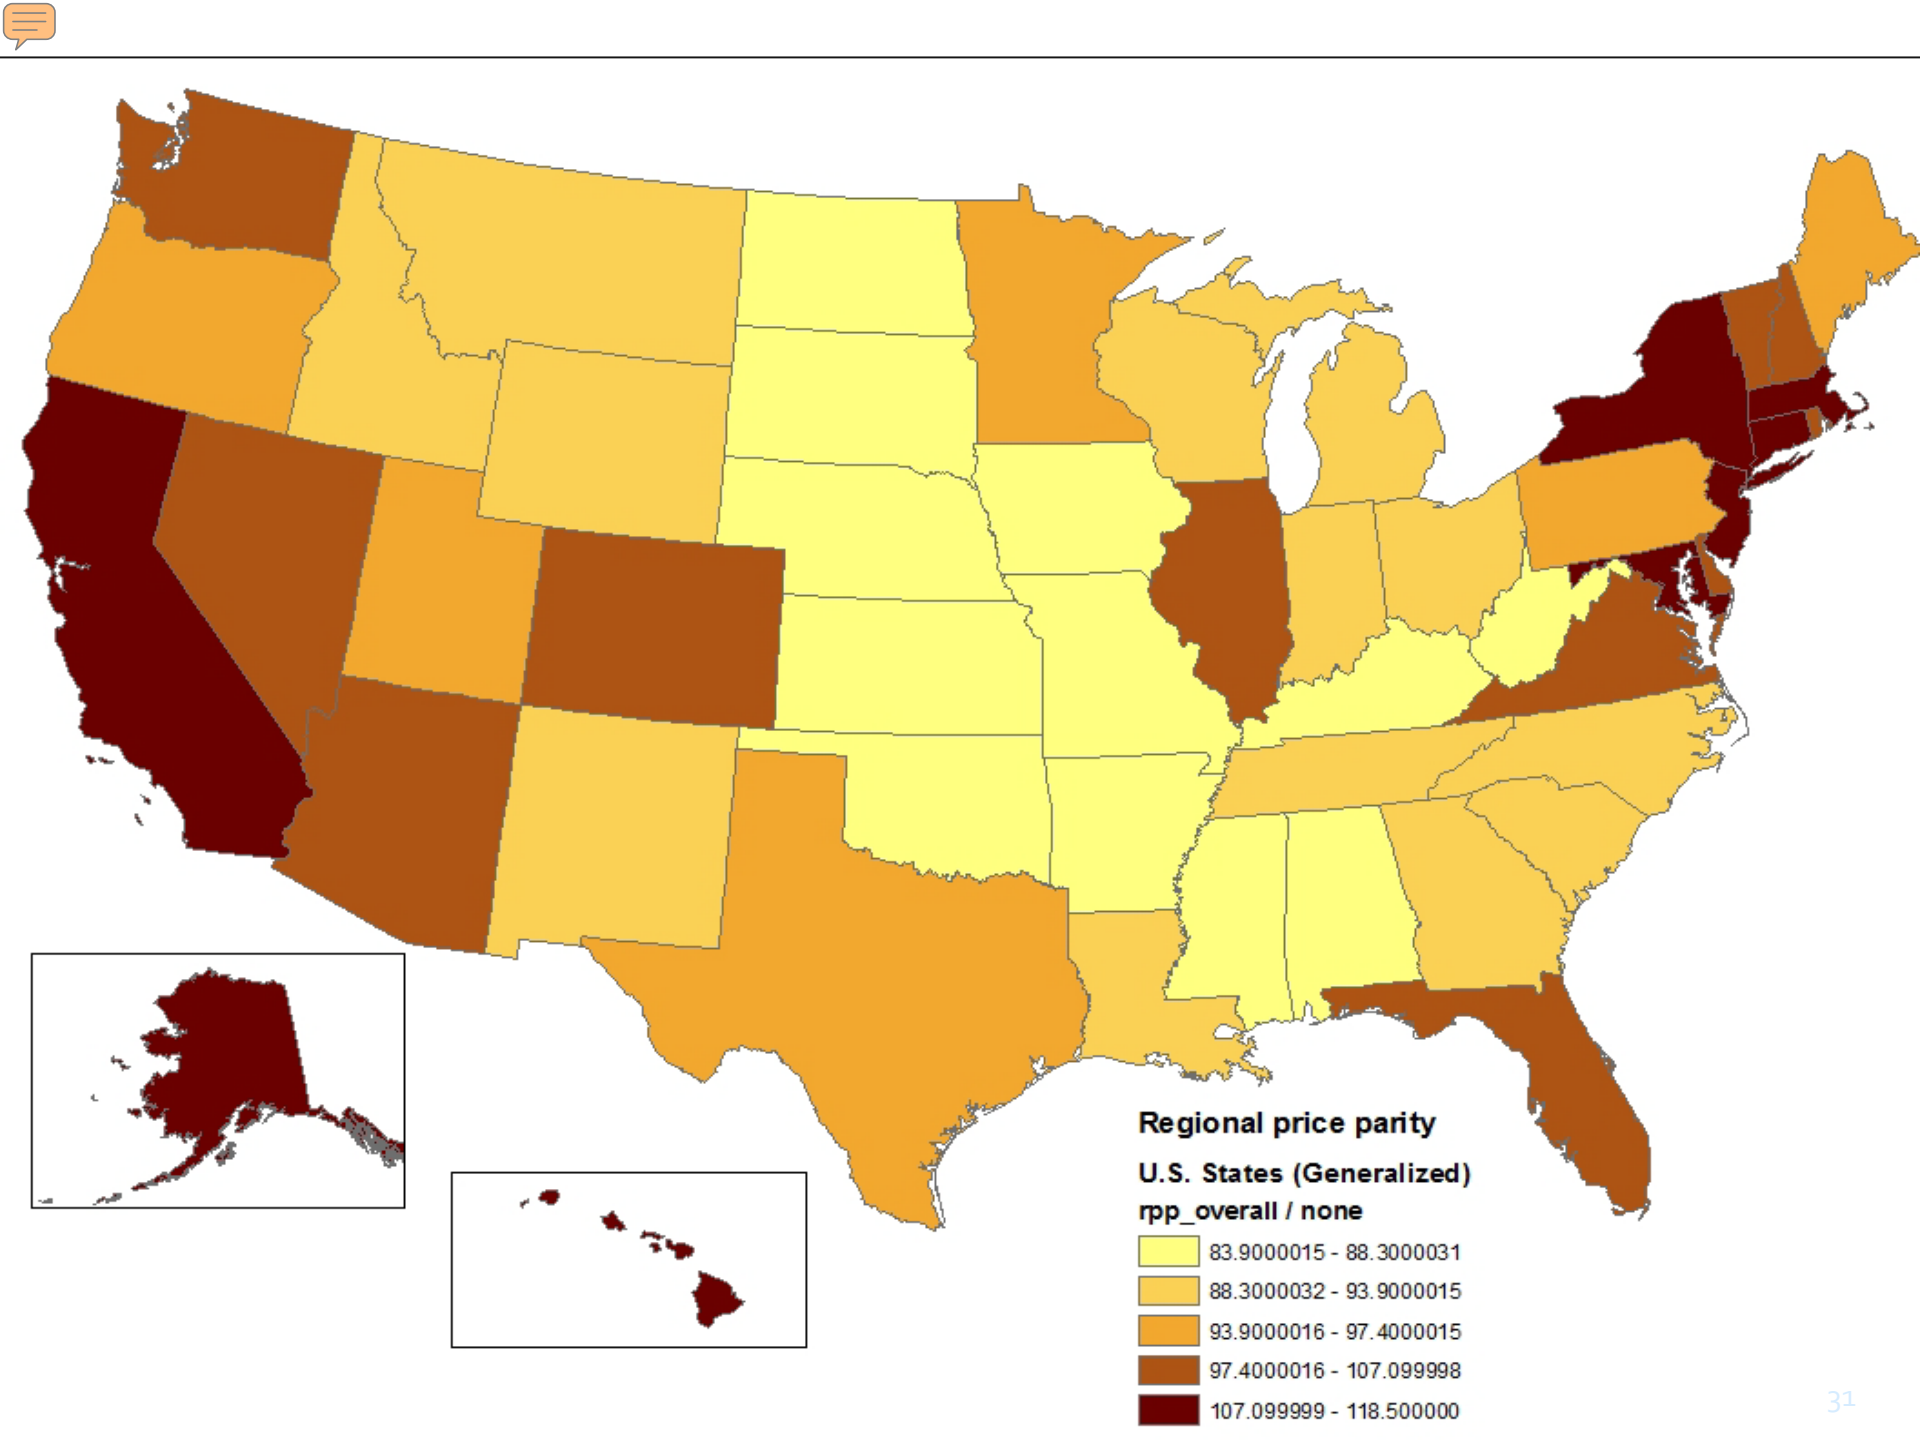
Regional price parity
U.S. States (Generalized)
rpp_overall / none
83.9000015 - 88.3000031
88.3000032 - 93.9000015
93.9000016 - 97.4000015
97.4000016 - 107.099998
107.099999 - 118.500000

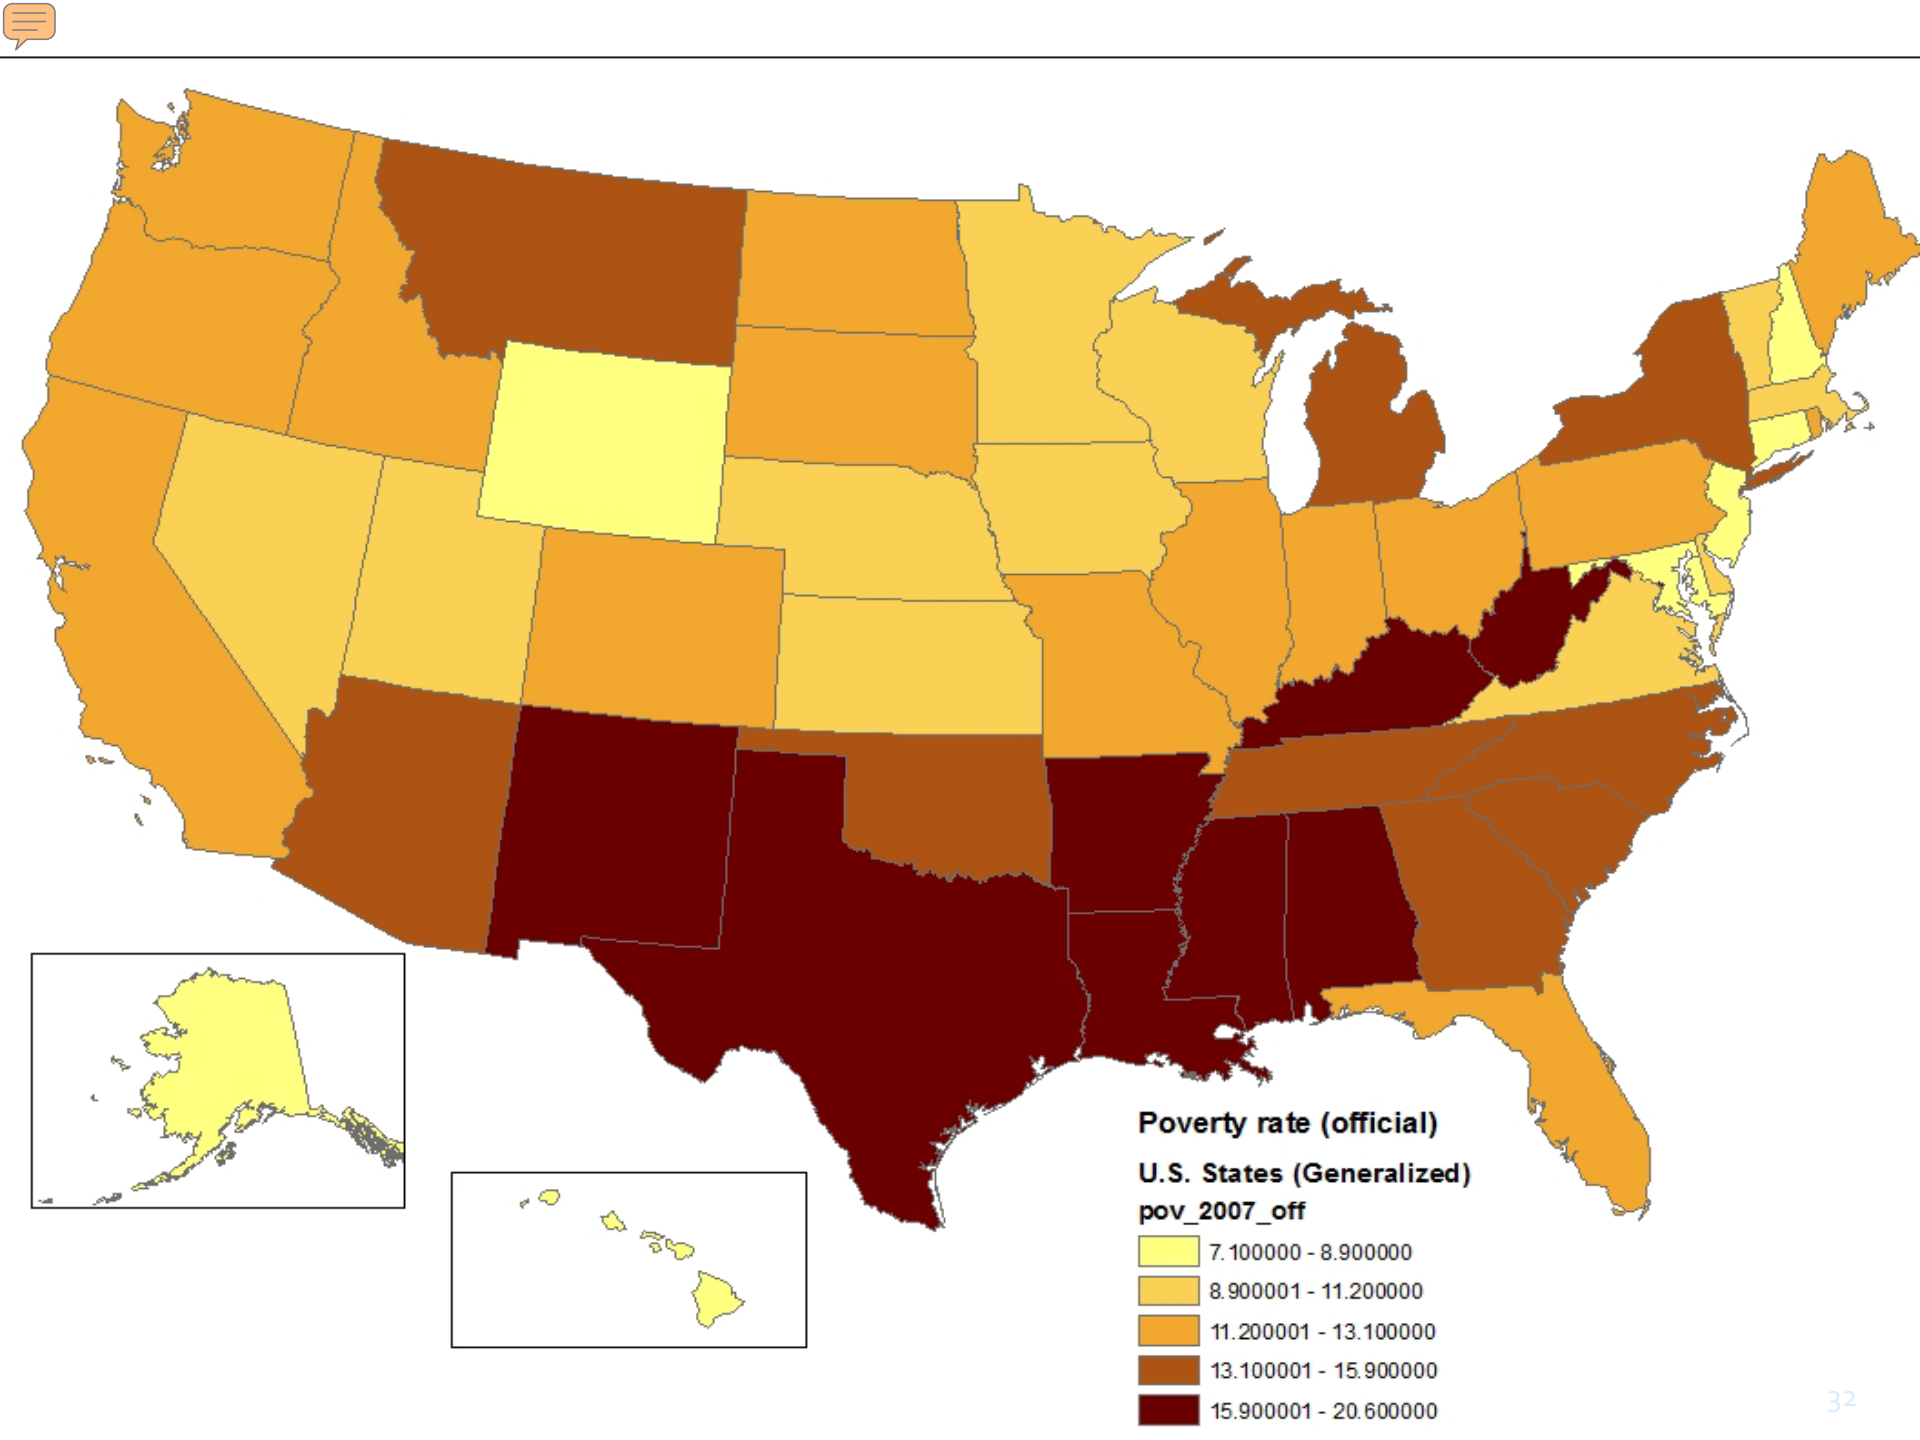
Poverty rate (official)
U.S. States (Generalized)
pov_2007_off
7.100000 - 8.900000
8.900001 - 11.200000
11.200001 - 13.100000
13.100001 - 15.900000
15.900001 - 20.600000

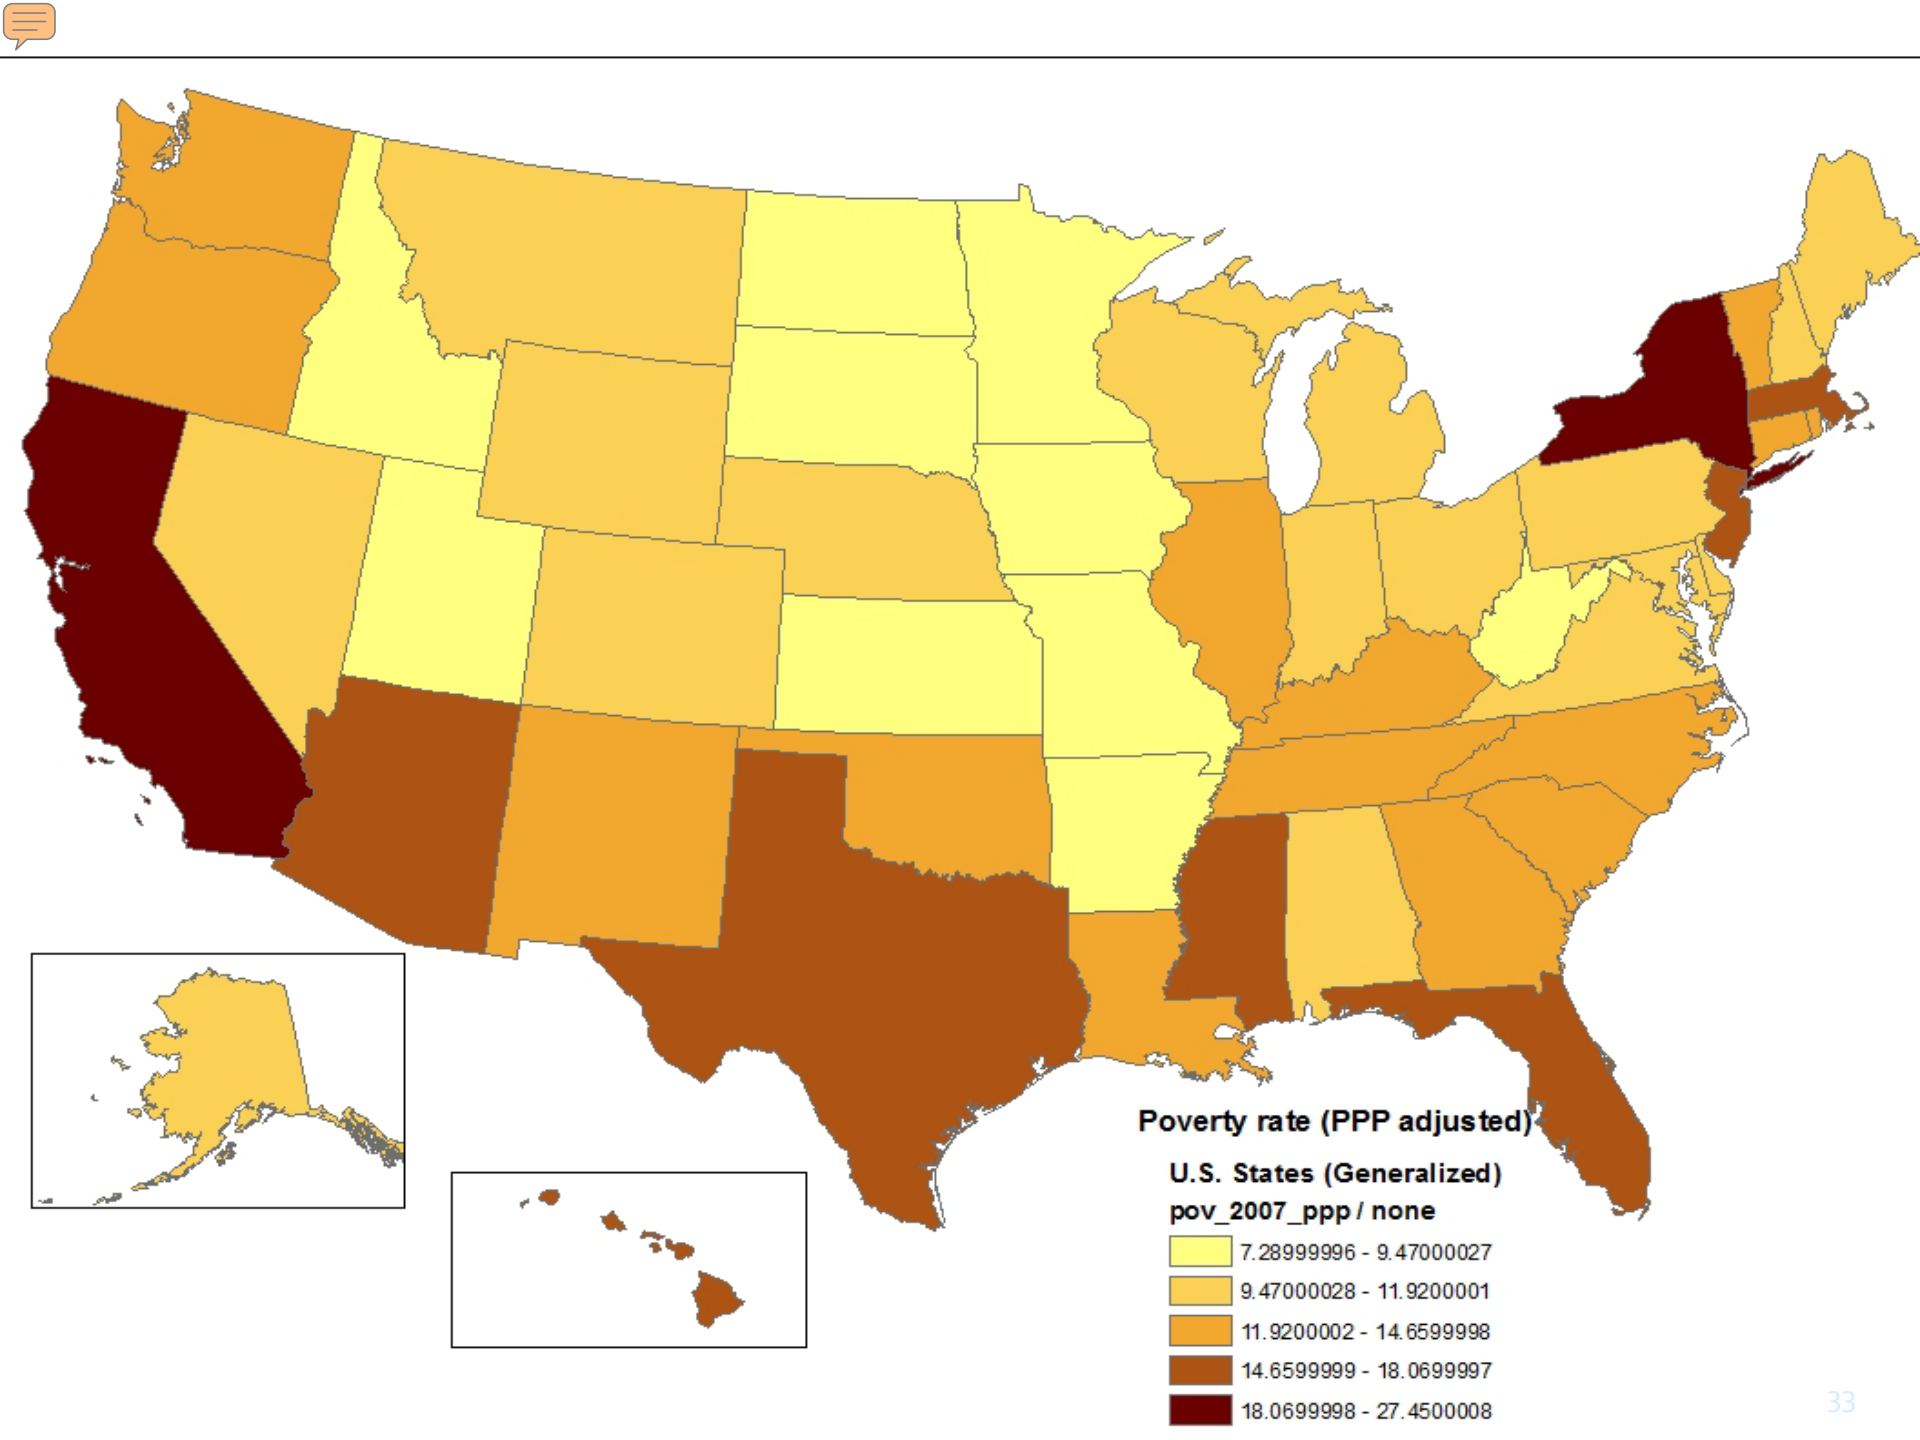
Poverty rate (PPP adjusted)
U.S. States (Generalized)
pov_2007_ppp / none
7.28999996 - 9.47000027
9.47000028 - 11.9200001
11.9200002 - 14.6599998
14.6599999 - 18.0699997
18.0699998 - 27.4500008

Correlations & regressions

- Across states, standard of living measure is correlated 0.43 with adjusted per capita income
  - Only 0.39 with unadjusted
- Ladder is correlated 0.37 with and 0.39 without
- Regression of mean ladder on lny and ln(RPP) not sig. on ln(RPP)
- Regression of SOL on lny and ln(RPP) have equal and opposite signs and are both (just) significant
- What people tell you doesn't do so badly
  - At least about standard of living
- Needs refined to lower income people

What do I take away?

- I am (mildly) in favor of geographical price adjustment in poverty measures
- We then adjust for *economic* well-being and ignore variations in amenities
  - Poor, but good amenities: can we handle that?
- Not trying to measure poverty in well-being space
- Very preliminary empirical work on well-being
- There are important amenities other than low prices
- Some evidence that prices are relevant for SWB
  - Are the variations in the Gallup data consistent with free mobility?
  - Yes, even if everyone is in the best place for them, there will be variations in outcomes across places
  - The people who live in Montana are *different* from those who live in NJ
- Asking people directly is worth pursuing further