Heterogeneity in the Impact of Economic Cycles and the Great Recession: Effects Within and Across the Income Distribution

Marianne Bitler
Department of Economics, UC Irvine and NBER
mbitler@uci.edu

Hilary Hoynes
Department of Economics and Public Policy, UC Berkeley and NBER
hoynes@berkeley.edu

January 12, 2015

ABSTRACT

In this paper, we examine the effects of economic cycles on low- to moderate-income families. We use variation across states and over time to estimate the effects of cycles on the distribution of income, using fine gradations of the household income-to-poverty ratio. We also explore how the effects of cycles affect the risk of falling into poverty across demographic groups, focusing on age, race/ethnicity and family type. We conclude by testing to see whether these relationships have changed in the Great Recession. We discuss the results in light of the changes in the social safety net in recent decades.

This paper was prepared for the AEA Annual Meeting in Boston January 2015 as part of the session "Impacts of the Great Recession on Low-Income Households". We thank Dorian Carloni and Elira Kuka for excellent research assistance and Diane Schanzenbach for helpful comments.

The Great Recession led to historic reductions in employment, earnings, and income for workers and families in the United States. The impacts of recessions are not shared equally across groups. Prior work has shown that the impacts are felt most strongly by men, black and Hispanic workers, youth, and low-education workers (for example see Hoynes, Miller and Schaller 2012). An analysis of workers, however, does not create a complete picture of the wellbeing of American families and how they are affected by economic cycles. Here, we extend the literature by estimating the effect of cycles both within and across the income distribution.

As with the prior labor market literature, we analyze the effects of cycles on individuals. However, we focus on the effects on household after-tax-and-transfer income (ATTI) (rather than individual employment or earnings). The advantage of using household ATTI is that it captures the cumulative effects of recessions on all of the (potential) workers and non-workers in the household and the effects of both the cash and noncash safety net and the tax system. We use household ATTI and observed household size to calculate household income-to-poverty ratios, using the official poverty thresholds associated with each household size. We then examine the effects of cycles across the income-to-poverty distribution, and can compare the effects of a one percentage point increase in the unemployment rate on the propensity to have income below 50% of the poverty threshold, 100% of the poverty threshold, 200% of the poverty threshold and so on.

We present three sets of results. First, we estimate the effect of cycles on the distribution of income-to-poverty ratios for the full sample and full period. Once we establish these basic findings, in our second set of results we examine how these effects vary for different demographic groups defined by age, race/ethnicity and family type. In our final set of results, we examine whether the impacts across the income-to-poverty distribution have changed between the early 1980s recession and the Great Recession.

We build our estimation sample using the Current Population Survey, providing annual data covering calendar years 1980-2013. We collapse the data to state by year (by demographic group) cells and estimate state panel data models controlling for state and year fixed effects. We measure the economic cycle using the state unemployment rate. The effect of cycles is identified using variation across states and over time in the timing and severity of recessions. We then extend our basic model to examine differences across groups and to see if the experience in the Great Recession differs from the historical relationship between unemployment and the income-to-poverty ratio.

This study builds on a rich and substantial literature that examines the connection between labor market opportunities, economic growth and poverty.[1] In our own recent work we developed our measure of ATTI poverty (Bitler and Hoynes 2010, 2014) and estimated the impact of cycles on nonelderly poverty (Bitler and Hoynes 2014) and child poverty (Bitler, Hoynes and Kuka 2014). Here we extend that work in several ways. First, we more comprehensively examine the effects across the income distribution, comparing the effects of cycles at different points of the income-to-poverty distribution. Second, we compare the effects across groups defined by age (children, prime aged adults, the elderly), race and ethnicity, and family type (families with married versus single heads).

Our analysis yields several important findings. We find that poverty rises in recessions and falls in expansions, and the level of cyclicality is substantially higher at the lowest levels of the income-to-poverty distribution. Elderly poverty rates are much less affected by cycles than are other age groups. On the other hand, for given levels of income-to-poverty, there is little difference in the cyclicality across other demographic groups (race/ethnicity, gender, family

[1] For example see Bitler and Hoynes 2010; Blank 1989, 1993; Blank and Blinder 1986; Blank and Card 1993; Cutler and Katz 1991; Freeman 2001; Gunderson and Ziliak 2004; Hines, Hoynes and Krueger 2001; Hoynes, Page and Stevens 2006; Meyer and Sullivan 2011.

structure). Finally, we find that the Great Recession has led to a significant increase in the cyclicality of non-elderly poverty at the lowest income-to-poverty levels (e.g., below 50% poverty) while the cyclicality has increased modestly at higher points of the income-to-poverty distribution (e.g., above 150% poverty).

I. Measuring Poverty

We use the Annual Social and Economic Supplement (ASEC) to the Current Population Survey (CPS), providing annual income data for the calendar years 1980-2013 (survey years 1981-2014).[2] The ASEC provides labor market outcomes, income, and program participation for the previous calendar year, as well as demographic information from the time of the survey. This survey is used to report official poverty each year in the United States (deNavas-Walt and Proctor 2014), and more recently the Supplemental Poverty Measure (Short 2014).

Based on the CPS, we construct household after-tax-and-transfer income (ATTI). Our measure of ATTI includes cash income plus the cash value (as reported by the household or imputed by the Census Bureau) of non-cash programs (food stamps, school lunch, housing subsidies, energy subsidies) and subtracts taxes (payroll tax, state taxes and federal taxes including the EITC, child and child care tax credits, and stimulus payments). The household's ATTI is the sum across all persons in the household (after dropping the very small number of "unrelated" children). We combine this with the official poverty thresholds (deNavas-Walt and Proctor, 2014) for each household (using observed household size and structure) to create the income-to-poverty ratio, which is assigned to everyone in the household. For more information on the construction of our ATTI poverty measure and how it differs from official poverty (and

[2] The CPS surveys for calendar years 1987 and 1991 are missing non-cash safety net programs; therefore our data are missing for those years.

the SPM) see Bitler and Hoynes (2014).[3]

II. Empirical Approach

We use a state panel data model, controlling for fixed state characteristics and common shocks where the economic cycle is measured by the state unemployment rate. In particular, we estimate:

$$(1) \qquad y_{st} = \beta UR_{st} + \alpha_s + \delta_t + \varepsilon_{st}$$

where y_{st} is the share of individuals in state *s* in year *t* who have ATTI below some multiple of the threshold (e.g., below 150% of poverty), UR_{st} is the state unemployment rate, and state and year fixed effects are given by α_s and δ_t respectively. We cluster the standard errors at the state level, and the regressions are weighted using the relevant denominator (the relevant CPS total weighted population in the state-year cell). Our empirical strategy exploits variation in the timing and severity of cycles across states to estimate the effect of labor market conditions on the income distribution.

We extend the basic model to explore whether the effect of the unemployment rate on the distribution of income-to-poverty in the Great Recession represents a change from historical patterns. We modify (1) and estimate:

$$(2) \qquad y_{st} = \beta_{80} D_{80} UR_{st} + \beta_{GR} D_{GR} UR_{st} + \beta_O D_O UR_{st} + \alpha_s + \delta_t + \varepsilon_{st}$$

We split 1980-2013 into three periods: the 1980s recessions and expansions ($D_{80} = 1$, 1980-1989), the Great Recession and expansion ($D_{GR} = 1$, 2007-2013) and the rest of period ($D_O = 1$, 1990-2006) and estimate the responsiveness (β_k) across the periods.

[3] The key differences are two-fold. First, official poverty uses cash, pre-tax income and second it measures income at the family (rather than household) level and compares it to the thresholds for the family (not household) size.

III. Impacts of Cycles within and Across the Income Distribution

We begin by estimating the effect of unemployment on income-to-poverty for all nonelderly persons (we turn to the elderly below). We estimate equation (1) for a series of outcome variables, the share of nonelderly persons in the state-year that have ATTI below 25% of the poverty threshold, 50% of the poverty threshold, 75% of the poverty threshold, and so on up to below 400% of the poverty threshold. To provide some context, for a family of three in 2013, the poverty threshold is $18,769, implying an income of $9,384 for 50% poverty, about $37,500 for 200% poverty, and about $75,000 for 400% poverty. Because of the significant variation in the fraction of families in each of these groups, we divide the coefficient in equation (1) by the (full period) mean of the dependent variable, and present percent impacts. The main results are presented in Figure 1 where we present the estimated percent impacts along with the 95 percent confidence intervals. For example, the coefficient plotted for 100% ATTI poverty implies that a one percentage point increase in the unemployment rate leads to a 6.2 percent increase in the share of persons with household ATTI below 100% poverty.

There are several findings evident from Figure 1. First, across the income-to-poverty distribution, the effects are positive and statistically significantly different from zero, showing that recessions (expansions) reduce (increase) incomes across the distribution.[4] Second, throughout most of the income-to-poverty distribution, the effects of unemployment on income-to-poverty are declining in income. That is, lower income-to-poverty levels are more affected by

[4] Note that we are estimating a number of models for different outcomes on the same data, which may raise concerns about multiple testing, and over-rejection. We adjust for this multiple testing using the Bonferroni correction to adjust the p-values to control the family-wise error rate (this controls the probability of falsely rejection one of a family of null hypotheses). In our context, this means a result is significant at the 5% level (using the Bonferroni adjustment) if 16 times the original p-value is less than .05, as there are 16 different outcomes. The Bonferroni correction is known to be conservative, despite this, we reject all 16 of these estimates being zero at the 5% level after adjusting the p-values in this fashion.

recessions than are higher income-to-poverty levels. This follows the results in the labor market literature whereby workers with lower education levels experience greater employment and earnings losses in recessions (e.g., Hoynes et al. 2006). Third, and importantly, the very lowest income levels, below 50% of poverty and below 25% of poverty, experience somewhat lower cyclical effects in percent terms compared to somewhat higher levels of income to poverty (below 75% or 100% of poverty).

It is important to note that the share of the distribution in each of these bins (0-25%, 25-50%, 50-75%, etc.) is not uniform. To illustrate this point, we also show on Figure 1 the (full sample) share of the population in each bin, for example the share below 25% of poverty, the share at or above 25% but below 50% of poverty, the share at or above 50 and below 75% of poverty, etc. This shows that the share with ATTI below 25% poverty is 0.014 and the share between 25 and 50% poverty is 0.013. Thus, the very large (percent) effects at the bottom of the distribution are affecting relatively few people.[5]

The results in Figure 1 refer to nonelderly persons. In Figure 2, we present results at the different points of the income-to-poverty distribution (share with ATTI below each cut of poverty) and examine how the effects vary across individuals of different age groups, broken into children (<18), prime age adults (18-64), and the elderly (65 and older). As with Figure 1, each point comes from estimation of equation (1), expressed as percent effects, along with the 95 percent confidence intervals. The effects of cycles on children and prime age adults are quite cyclical at low cuts of income to poverty. By marked contrast, elderly individuals have small and statistically insignificant effects of cycles on ATTI below around 125% poverty, and are much

[5] Note that the regression estimates refer to cuts on the income-to-poverty distribution that are cumulative – ATTI below 25% poverty, below 50% poverty, etc. The histogram shares, though, are not cumulative but correspond to the bins (<25%, 25-50%, 50-75%, etc).

smaller than those for the non-elderly until relatively high up the income distribution.[6] This may be due to the elderly experiencing less exposure to the labor market and more protection through Social Security. We have also broken the estimates down by race/ethnicity and by marital status and gender of the family head. There is strikingly little difference in the cyclicality of income-to-poverty across the race and ethnicity of the family heads. This implies that the greater cyclicality of average income for Hispanics and African Americans (compared to whites) is explained by differences in the distribution of incomes across these groups (rather than group per se). There is somewhat more of a difference across marital status groups, with unmarried female heads at the bottom of the income-to-poverty distribution being somewhat less responsive to the unemployment rate.[7]

Next, we turn to comparisons across the two most recent deep recessions, the Great Recession and the two consecutive severe downturns of the early 1980s. We estimate equation (2), generalizing our model to allow for differential effects during the Great Recession (2007-2013), the early 1980s period (1980-1989) and the rest of the period (1990-2006). As with Figure 1, we estimate a series of models with different cuts across the income-to-poverty distribution. We plot percent effects and 95% conference intervals, and compare the effects in the Great Recession to the 1980s (we consider the rest of period to be incidental parameters, but these results are available on request).[8] We show this broken down by age group, with Figure 3a showing the coefficients for the Great Recession and 1980s recessions for children and Figure 3b showing the coefficients for the same two periods for those 18-64. In each figure, the estimates

[6] Again, within each subgroup, we adjust the p-values using the Bonferroni correction, known to be conservative. 13 of the 16 estimates for children retain significance at the 5% level (the non-significant values include being below 25% of poverty, 375% of poverty, and 400% of poverty).. All of the values are significant for the prime aged adults, and none are for the elderly.

[7] Due to space limitations, these figures are not included here but are available on request.

[8] We use the full period mean for constructing the percent effects and therefore we plot $\beta_{GR}/\bar{y}$ and $\beta_{80}/\bar{y}$.

for the Great Recession are given by blue squares and for the 1980s recessions they are given by black circles. The square symbols for Great Recession are filled if the coefficient is statistically significantly different from the 1980s coefficient at the 5 percent level.

There are several important findings in Figures 3a and 3b. The most striking differences in the effect of unemployment on poverty are at the lowest levels of the income-to-poverty ratio. While the 1980s period shows an overall "inverted U shaped pattern" with smaller percent effects at the lowest and highest income-to-poverty levels, the Great Recession shows the largest effects at the lowest income-to-poverty levels and instead a steep gradient, with cyclicality declining as we move up the income distribution. For adults, the higher cyclicality during the Great Recession extends throughout the entire income-to-poverty distribution. For children, 75% and 100% poverty show substantially lower cyclicality in the Great Recession, but at 150% of poverty and higher, incomes are more cyclical. Overall, though, few of the estimates are significantly different for the children (<100%, <275%, <300%), but they are significantly different for adults at the very bottom (<25% and <50%) as well as at <175% and above. Finally, although not shown here, there is little difference in responsiveness across these two periods for the elderly.[9]

In earlier work, we have argued that changes in the social safety net may explain some of these differences over time (Bitler and Hoynes 2014). The increase in cyclical sensitivity at the bottom of the distribution seems clearly tied to welfare reform and the dramatic decrease in welfare caseloads and take-up of TANF (changes which should only have affected households with children). In contrast to the pre-welfare reform era, in the post welfare reform era TANF caseloads and expenditures have no relationship to "need", as measures by state labor market

[9] Adjusted these for the multiple hypotheses using the conservative Bonferroni correction, none of the differences remain significant for children, and only that for being under 275% of poverty is significant for the prime age adults.

conditions (Bitler and Hoynes 2014). Because the generosity of AFDC never reached much above extreme poverty, the increased exposure to the cycle is concentrated at the very bottom of the distribution. We also find somewhat greater protection in the Great Recession due to SNAP, which could explain the lower cyclicality for children at 75% and 100% of poverty. The greater cyclicality in the Great Recession experienced by income-to-poverty between 125% poverty to 400% poverty is likely due to the changes in the composition of the unemployed, with historic increase in long term unemployment (Valetta 2013) which we are not accounting for in our analysis.

V. Conclusion

In this paper, we have comprehensively examined the effects of the business cycle on the distribution of income-to-poverty. We find that effects are more cyclical at very low levels of income-to-poverty, and then become increasingly smaller as one goes up the income distribution. This gradient has become steeper in the Great Recession (when compared to the early 1980s recession), with the most disadvantaged being relatively more affected (compared to higher income levels) in the Great Recession. We also have explored differences across groups, finding large distinctions between the cyclical effects for both children and prime age adults and the correspondingly lacking evidence of cyclicality for the elderly. There were few significant differences across other demographic groups such as race/ethnicity, gender or marital status of the family heads.

REFERENCES

Bitler, Marianne and Hilary Hoynes. (2010). "The State of the Safety Net in the Post-Welfare Reform Era," Brookings Papers on Economic Activity: 71-127.

Bitler, Marianne and Hilary Hoynes. (2014). "The More Things Change the More They Stay the Same? The Safety Net and Poverty in the Great Recession," Forthcoming, Journal of Labor Economics.

Bitler, Marianne, Hilary Hoynes, and Elira Kuka. (2014). "Child Poverty and the Great Recession in the United States," working paper.

Blank, Rebecca. (1989). "Disaggregating the Effect of the Business Cycle on the Distribution of Income," Economica, 56:141-163.

Blank, Rebecca. (1993). "Why Were Poverty Rates so High in the 1980s?" In Poverty and Prosperity in the Late Twentieth Century. Papadimitriou and Wolff, eds., St Martin's Press.

Blank, Rebecca and Alan Blinder. (1986). "Macroeconomics, Income Distribution, and Poverty," In, Sheldon Danziger and Daniel Weinberg, eds., Fighting Poverty: What Works and What Doesn't. Harvard University Press.

Blank, Rebecca and David Card. (1993). "Poverty, Income Distribution and Growth: Are They Still Related?" Brookings Papers on Economic Activity 2.

Cutler, David and Larry Katz. (1991). "Macroeconomic Performance and the Disadvantaged," Brookings Papers on Economic Activity 1.

DeNavas-Walt, Carmen and Bernadette Proctor (2014). "Income and Poverty in the United States: 2013." Current Population Reports P60-251. U.S. Census Bureau.

Freeman, Richard. (2001). "The Rising Tide Lifts…?" In, Sheldon Danziger and Robert Haveman, eds., Understanding Poverty. Cambridge, MA: Harvard University Press, 97-126.

Gundersen, Craig and James Ziliak. (2004). "Poverty and Macroeconomic Performance across Space, Race, and Family Structure," Demography, 41:1, 61-86.

Hoynes, Hilary W., Douglas Miller, and Jessamyn Schaller. (2012). "Who Suffers in Recessions and Jobless Recoveries?" Journal of Economic Perspectives 26(3), 27-48.

Hoynes, Hilary, Marianne Page and Ann Huff Stevens (2006). "Poverty in America: Trends and Explanations" Journal of Economic Perspectives, 20(1):47-68.

Hines, James, Hilary Hoynes and Alan Krueger (2001). "Another Look at Whether a Rising Tide Lifts all Boats," in The Roaring Nineties: Can Full Employment Be Sustained? edited by Alan Krueger and Robert Solow. Russell Sage Foundation: New York.

Meyer, Bruce and James Sullivan. (2011). "Consumption and Income Poverty over the Business Cycle," Research in Labor Economics, vol. 32, 51-82.

Short, Kathleen. (2014). "The Research Supplemental Poverty Measure: 2013." Current Population Reports P60-251. U.S. Census Bureau.

Valetta, Robert (2013). “Long-term Unemployment: What Do We Know?” Economic Letter, Federal Reserve Bank of San Francisco.

Figure 1: Percent impacts of the unemployment rate on the household after tax and transfer income to poverty ratio, non-elderly population

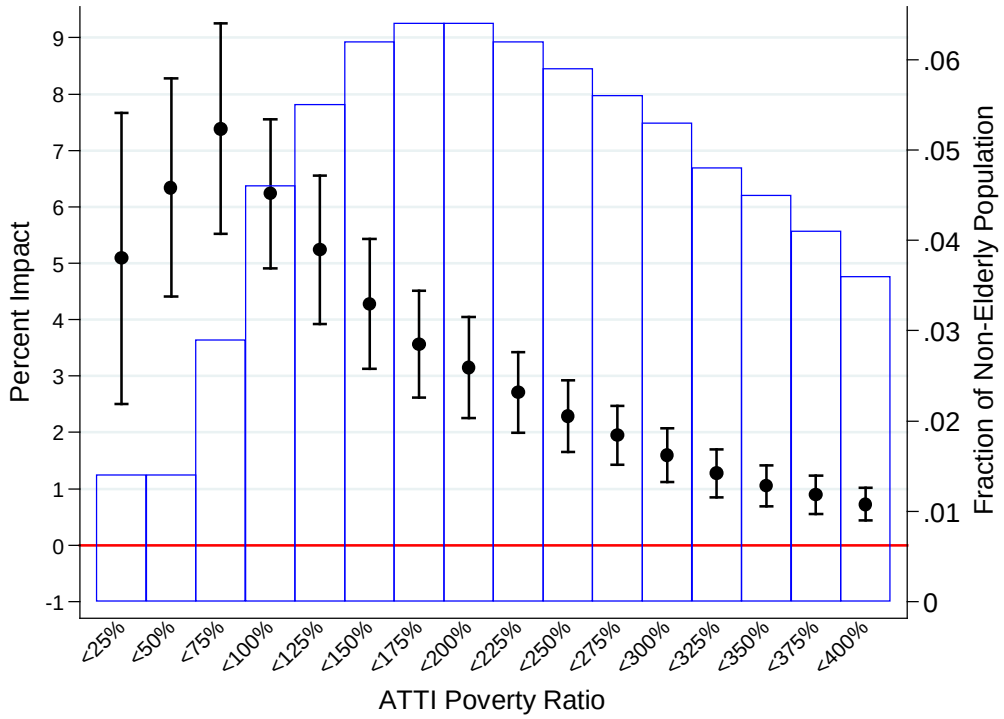

Notes: The figure presents (on the left scale) the coefficient estimate (and 95% CIs) for the effect of a one percentage point increase in the unemployment rate on the propensity to be below various cuts of the household income-to-poverty ratio, estimated from a series of regressions that also control for state and year fixed effects, for the nonelderly population. Household income-to-poverty ratio is calculated using ATTI income and household number of persons compared to the official poverty thresholds, averaged for each state-year cell using the CPS ASEC for calendar years 1980-2013. Regressions are weighted using the sum of person weights within each state-year cell, and are clustered at the state level. Percent effects divide the coefficient on the unemployment rate by the mean of the dependent variable (mean is taken over the entire period). The graph also contains a histogram (on the right scale) for the share of the non-elderly population that is between adjacent cuts of the income to poverty ratio. 25 percent of the population is above 400% of the household after tax and transfer income to poverty ratio.

Figure 2: Percent impacts of the unemployment rate on the household after tax and transfer income-to-poverty ratio, by age group

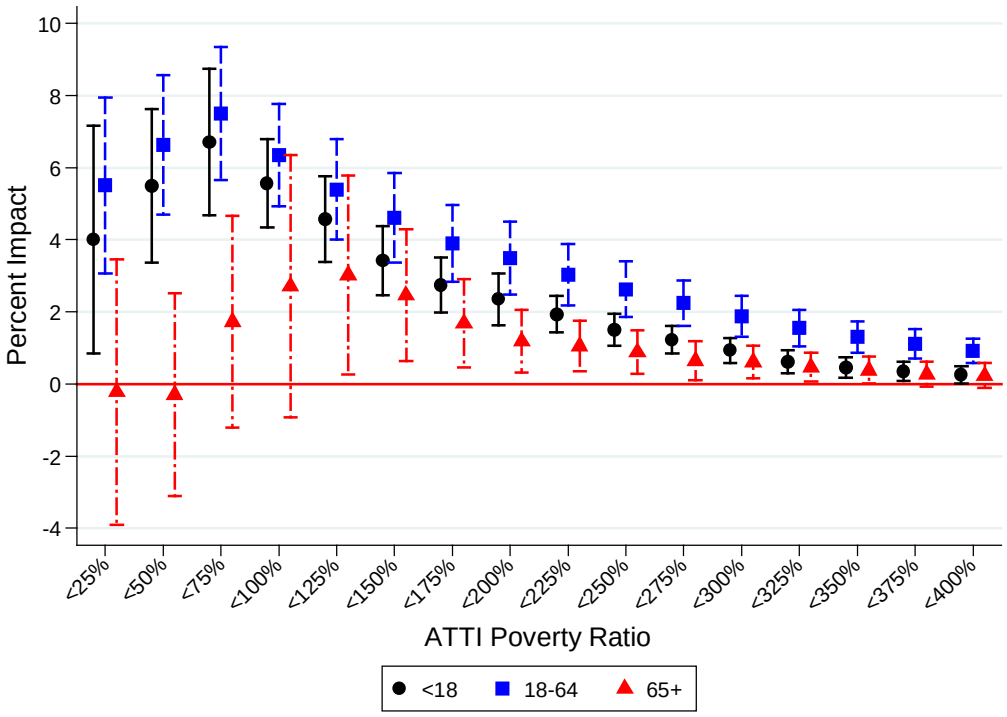

Notes: The figure presents the coefficient estimates (and 95% CIs) for the effect of a one percentage point increase in the unemployment rate on the propensity to be below various cuts of the household income-to-poverty ratio, by age group. The black circles are the estimates for children (<18), the blue squares are for those 18-64, and the red triangles are for the elderly. Each regression also controls for state and year fixed effects. Household income-to-poverty is calculated using ATTI income and household number of persons compared to the official poverty thresholds, averaged for each state-year cell using the CPS ASEC for calendar years 1980-2013. Regressions are weighted using the sum of person weights within each state-year cell, and are clustered at the state level. Percent effects divide the coefficient on the unemployment rate by the mean of the dependent variable (mean is taken over the entire period).

Figure 3: Percent impacts of the unemployment rate on the household after tax and transfer income-to-poverty ratio, Great Recession versus 1980s recessions

(a) Ages 0-17

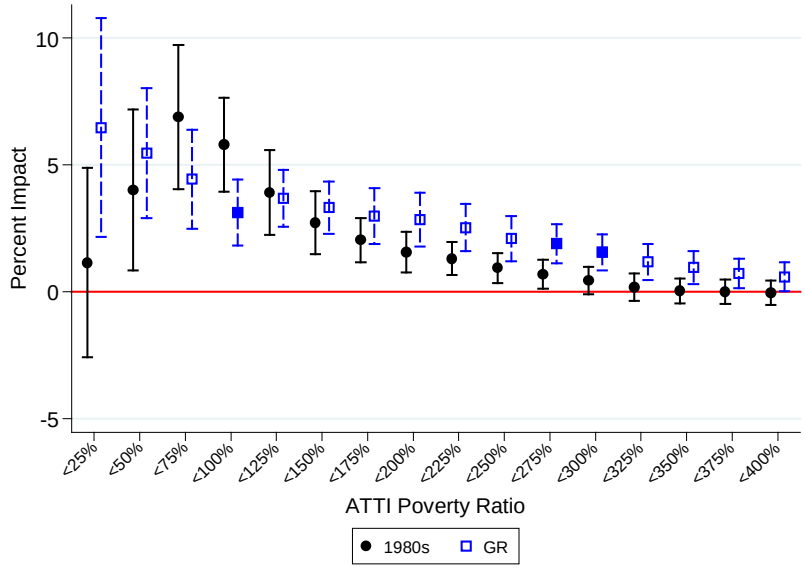

(b) Ages 18-64

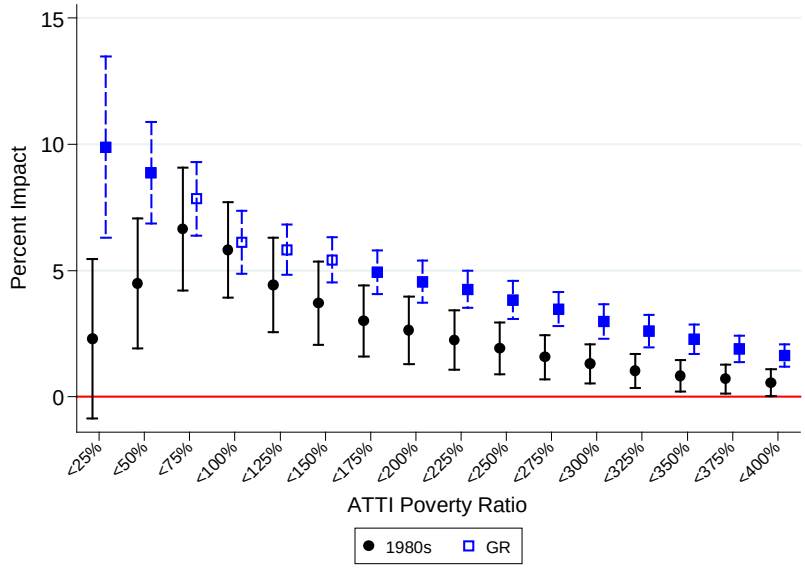

Notes: Each figure presents the coefficient estimates (and 95% CIs) for the effect of a one percentage point increase in the unemployment rate on the propensity to be below various cuts of household income-to-poverty, for the Great Recession (blue squares) and 1980s recessions (black circles). Square symbols for Great Recession are filled if the coefficient is statistically significantly different from the 1980s coefficient at the 5 percent level. Estimates for Figure 3a are for children and those for Figure 3b are for 18-64 year olds. Each regression also controls for state and year fixed effects. Household income-to-poverty is calculated using ATTI income and household number of persons compared to the official poverty thresholds, averaged for each state-year cell using the CPS ASEC for calendar years 1980-2013. Regressions are weighted using the sum of person weights within each state-year cell, and are clustered at the state level. Percent effects divide the coefficient on the unemployment rate by the mean of the dependent variable (mean is taken over the entire period).