# Racial Labels and the Measurement of Racial Inequality: The Case of Concentrated Poverty

Dionissi Aliprantis[+] David Van Riper[*]

September 12, 2026

**Abstract:** Large efforts in research and policy are driven by the fact that in segregated American cities there is a high concentration of poor Black residents in neighborhoods with the lowest socioeconomic status (SES). This concentration has experienced a massive decline over the past decade as measured in the American Community Survey (ACS). This decline is not due to tract boundary changes. This decline could be due to changes in the neighborhood SES of poor Black Americans, but we show that in most cities we study, a large share could instead be due to changes in the way the Census Bureau assigns racial labels to individuals. This example highlights how changing definitions of race complicate measuring trends in racial inequality.

**Keywords:** Concentrated poverty, residential segregation by race, racial categories
**JEL Classification Codes:** C81, J11, J15, R23

[+]: ENS de Lyon, *dionissi.aliprantis@ens-lyon.fr*
[*]: Institute for Social Research and Data Innovation, University of Minnesota, *vanriper@umn.edu*
*Acknowledgements:* We thank Tim Dunne and Dylan Small for helpful comments. An earlier draft of this note was circulated under the title "The massive recent decline in concentrated poverty: A change of neighborhoods or of racial labels?"

# 1 Introduction

What definition of race should be used to describe the population of the United States (US)? Contemporary definitions of race have their origins in slavery. As Davis (2006) notes in his history of the Atlantic slave trade, the invention of race solved an old problem in the history of slavery.

> long before the eighteenth-century invention of 'race' as a way of classifying humankind, a different phenotype or physical appearance made the dehumanization of enslavement much easier. ... Throughout the ancient Euro-Asian world as well as in the preconquest Western Hemisphere, slaves were commonly marked off by identifying symbols or icons, such as brands, tattoos, collars, hairstyles, or clothing. Clearly such emblems would have been less necessary if all slaves had shared distinctive physical characteristics that quickly differentiated them from all nonslaves (p. 53).

After slavery ended, a primary use of race remained the justification of limiting either the rights or access to resources of those in the out-group (Zuberi (2001), Washington (2006), Mason (2023)).

Contemporary social scientists typically conduct research on racial inequality in terms of the groups defined in 1997 by the US Office of Management and Budget (OMB (1997)). Given the history described above, alternative definitions of racial groups should be a topic of discussion for those wanting the US to be a society with equal rights and equal opportunity. It is therefore laudable that the US Census Bureau has conducted extensive research and outreach to improve data on race and ethnicity (US Census Bureau (2024), Jones et al. (2024)). Nevertheless, while there are reasons that changes to the definition of race could be socially beneficial, accurately describing changes to society across such a discontinuity would require careful attention on the part of researchers and data users.

Starting with the 2020 decennial Census and American Community Survey (ACS), the US Census Bureau changed the way that it processes race and ethnicity data (See Appendix A for a full description.). As noted in Arias et al. (2025), researchers and data users have been provided limited guidance about these changes, which stands in contrast to the guidance given around the adoption of the 1997 OMB definition. In our research on residential segregation, we noticed striking changes around 2020. We were not sure whether to attribute those changes to the residential sorting of people across neighborhoods or to the way people are classified into race and ethnicity groups. Given the widespread effort of research and policy to understand and mitigate the effects of concentrated poverty, starting at least with the seminal work in Wilson (1987), an accurate interpretation of the data is a fundamental requirement for researchers and policy makers.

This note shows that changes to the US Census Bureau's measurement of race make it difficult to interpret the recent decline in concentrated poverty. We first show that this decline has been massive and widespread across racially-segregated cities. The magnitude of this change alone gives us pause as to whether it is an artifact of some detail in its measurement, even if overall poverty declined substantially over the same time period. We show that the decline in concentrated poverty is not due to the 2020 change in Census tract boundaries. While changes to tract boundaries is an obstacle to measurement in other contexts (Glaeser et al. (2025)), we find that the decline in concentrated poverty is even larger when using measures that account for tract boundary changes.

We show that the stability of labeling conventions is critical for interpreting the recent decline in concentrated poverty. We first document two secular trends that yield opposite perspectives on the stability of labeling conventions. There is a secular decline in the number of poor Black residents in low-SES neighborhoods that is stable across the 2020 measurement change. However, the secular trend in the number of poor residents in low-SES neighborhoods who are labeled as "Two or More Races" is not stable across the 2020 measurement change. In most cities, the increase in the number of "Two or More Races" individuals who are either "Black and White" or "Black and American Indian / Alaska Native" could explain the majority of the decline in poor Black residents in low-SES neighborhoods.

We conclude by showing how uncertainty about the stability of labeling conventions translates into uncertainty about secular changes in concentrated poverty. We calculate changes in concentrated poverty for the 30 most-segregated large cities in the US under two bounding assumptions about labeling conventions: (i) that labeling has been stable, and (ii) that all new "Two or More Races" individuals post-2020 would have been labeled as "Black" pre-2020. In one third of cities, uncertainty between assumptions (i) and (ii) makes no difference: the recent decline in measured concentrated poverty reflects genuine improvements in the neighborhood SES experienced by poor Black Americans. In the remaining two thirds, though, using assumption (ii) instead of assumption (i) changes the measured decline in concentrated poverty by an economically meaningful amount. For one third of cities the change is moderate, and for the final third the change is enough to entirely overturn the measured decline. In other words, in two thirds of large cities it is possible that a significant share of poor residents labeled as Black pre-2020 have experienced no change in neighborhood SES, but have simply been labeled with a different racial category post-2020.

It is worth emphasizing that the measurement of concentrated poverty studied in this note is simply one example highlighting a broader point: Secular changes in racial inequality are a combination of changes in outcomes and changes in the way racial labels are applied to individuals, both by society and by those individuals themselves (Saperstein and Penner (2012)).[1] To illustrate that this problem is not specific to the Census and ACS, consider another prominent survey, the National Longitudinal Survey of Youth 1979 (NLSY79). In early waves of the NLSY79, race was recorded in separate variables as reported by survey enumerators and as self-reported. These variables followed the US government's Office of Management and Budget (OMB)'s 1977 standards, which permitted only one racial label per person. Then, in 2002, respondents self-reported their race in questions conforming to the OMB's 1997 standards, which allowed for more than one racial label per person (Light and Nandi (2007)). Which of these measures a researcher uses is not innocuous; they generate economically significant differences in basic statistics like the unemployment rate of Black men (Williams-Baron and Saperstein (2026)). Our results on concentrated poverty represent an important example of how changing definitions of race create uncertainty in the measurement of racial inequality.

[1] The same is true for ethnic inequality; see Duncan and Trejo (2011) for analogous evidence.

# 2 The Recent Decline in Concentrated Poverty

Concentrated poverty is high in many cities in the United States (US). Consider the distribution of poor Black residents in a city by their neighborhoods' socioeconomic status (SES), where we construct SES as a ranking of Census tracts on a scale of 0 (lowest) to 100 (highest) in terms of poverty rate, high school diploma attainment rate, bachelor's degree attainment rate, the employment to population ratio, the unemployment rate, and the share of households with children under 18 that are single-headed.[2] As an example, the light blue bars in Figure 1a show that 33 percent of poor Black residents of Baltimore live in Census tracts ranked in the bottom five percent nationwide in terms of neighborhood SES. Furthermore, 42 percent of this group lives in the bottom 10 percent of neighborhoods. The red bars in Figure 1a show that poor white residents of Baltimore are uniformly distributed across neighborhood SES.[3]

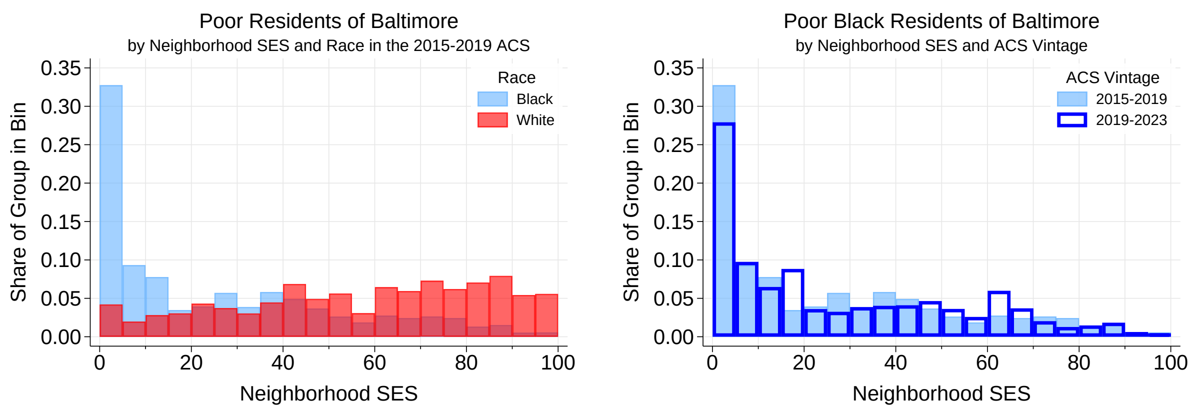


(a) Neighborhood SES and Race

(b) Change in Segregation by ACS Vintage

Figure 1: Racial Segregation in Baltimore,
by Vintage of the American Community Survey (ACS)

Note: These figures show the distribution of the Baltimore metro area's poor population over Census tracts ranked by a neighborhood socioeconomic status (SES) ranking whose construction we discuss in the main text. The left panel show the distributions conditional on a poor resident being either Black or non-Hispanic white. The right panel shows the distributions for poor Black residents in the 2015-2019 or 2019-2023 vintages of the American Community Survey (ACS).

[2]We obtain Census data from the National Historical Geographic Information System (Manson et al. (2024)). For the sake of this note, we will refer to concentrated poverty as the concentration of poor Black residents in neighborhoods with socioeconomic status (SES) in the bottom 10 percent of all neighborhoods in the US, as shown in Figure 1a. See Aliprantis et al. (2024) for a more detailed discussion of the construction of the neighborhood SES ranking used in this note and for a comparison of the strengths and weaknesses of this ranking with those from the Childhood Opportunity Index (COI, Noelke et al. (2020)) and the Opportunity Atlas (OA, Chetty et al. (2020)). We measure poverty using data from ACS' Table C17002, which reports the population distribution across ratio-of-income-to-poverty-threshold categories, based on the federal Office of Management and Budget's official poverty definition (Statistical Policy Directive 14).

[3]We focus on Baltimore due to the prominence of the Baltimore Regional Housing Partnership (BRHP) in the literature on housing mobility programs (Aliprantis and DeLuca (2025), DeLuca and Rosenblatt (2017), Darrah and DeLuca (2014), Aliprantis et al. (2024)).

Given the widespread concern in the social sciences about the effects of concentrated poverty since Wilson (1987), it is noteworthy that cities like Baltimore experienced a massive decline in concentrated poverty from the 2015-2019 American Community Survey (ACS) to the most recent 5-year vintage, 2019-2023. Looking at Figure 1b, the difference between the left-most solid light blue bars and the left-most outlined blue bars indicates that 4.4 percentage points fewer of Baltimore's poor Black residents live in neighborhoods ranked in the bottom 10 percent of neighborhood SES.

The speed of change in Baltimore is extremely high for a four-year period. If change were to continue at this pace, the segregation of poor Black residents of Baltimore would be resolved in about two decades. In comparison, it might take a housing mobility program 10 decades to resolve the segregation of poor Black residents in Baltimore (Aliprantis et al. (2024)).[4]

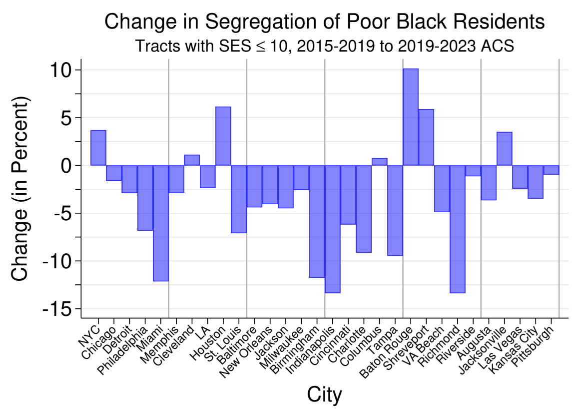


Figure 2: Change in Racial Segregation, by City

Note: This figure shows, for each of our sample cities, the change in the percent of poor Black residents living in the bottom 10 percent of tracts nationwide between the 2015-2019 and 2019-2023 vintages of the American Community Survey (ACS).

It is additionally surprising that the change measured in Baltimore is typical. We use a sample of large, highly-segregated cities in our analysis, defined as Core-Based Statistical Areas (CBSAs) with 50 thousand poor Black residents and at least 30 percent of those poor Black residents residing in tracts in the lowest 10 percent of neighborhood SES nationwide in the 2015-2019 ACS. Figure 2 shows the change in concentrated poverty for this sample of cities between the 2015-2019 and 2019-2023 vintages of the ACS. In this sample of cities, the median decline in concentrated poverty between the 2015-2019 and 2019-2023 ACS vintages is 3.2 percent. It is reassuring to establish the widespread nature of this phenomenon due to the difficulty of constructing standard errors for such a statistic.

For the sake of exposition, in the subsequent analysis will often refer to 5-year estimates of the American Community Survey (ACS) using their middle year. Since changes starting with the 2020 Census will first appear in the 2016-2020 ACS, this means we will often focus on changes after the

[4]Aliprantis et al. (2024) estimate that a full-blown housing mobility program would resolve about 10 percent of this type of concentrated poverty. Since one would imagine such a result over the course of approximately a decade, fully addressing the problem with this policy might take a century.

2017 (or 2015-2019) ACS and beginning with the 2018 (or 2016-2020) ACS.

# 3 Potential Explanations

## 3.1 Changes between 2010 and 2020 Census Tract Boundaries

We test whether Census tract boundary changes in 2020 are driving the measured declines in concentrated poverty using two measures of population counts that are robust to tract boundary changes. At each decennial Census, tract boundaries are updated so that tract population counts generally fall within the Census Bureau's expected range of 4,000 to 8,000 residents. Glaeser et al. (2025) show that these boundary changes can create major obstacles to measuring changes in population. Full details about the methodologies used are presented in Appendix E.

We first follow Glaeser et al. (2025) and create population counts area-weighted back into fixed 2010 tract boundaries. The 2015-2019 ACS is already recorded on 2010-vintage tract boundaries and requires no adjustment. The 2019-2023 ACS is reported for 2020-vintage boundaries, though. For this latter vintage we reallocate each 2020-boundary tract's poor Black population onto every 2010-boundary target tract it overlaps, using each 2020 tract's share of area falling in that target tract. We do this using the Census Bureau's official 2010-to-2020 Census Tract Relationship File, which records the area of geographic overlap between every pair of overlapping 2010- and 2020-vintage tracts. We similarly area-weight-average each source tract's neighborhood SES percentile onto the target tract (weighted by the target tract's own area share).

We also consider a second measure robust to boundary changes, Target-Density-Weighting (TDW; Schroeder (2007)). TDW addresses area-weighting's implicit assumption that population is spread uniformly across a tract's entire area, including any parks, water, or non-residential land it contains. TDW instead weights each tract-to-tract overlap by its area *and* by how population-dense the target tract is, so that overlaps landing in denser tracts receive more of the source tract's population than equally-sized overlaps landing in less-dense ones. We measure each target tract's density using the NHGIS's published crosswalk between 2020-vintage and 2010-vintage Census tracts (Schroeder et al. (2026)).[5] For every overlapping source-target tract pair, this crosswalk reports NHGIS's expected share of the source tract's total population located in the target tract. We apply this population share directly to each source tract's poor Black population and neighborhood SES quality, so our reallocated values reflect how NHGIS itself expects a source tract's population to be distributed across its target tracts. Appendix E also reports a third robust measure, which uses a TDW approach but instead calculates each target tract's density from actual 2020 Census block-level Black population.

Accounting for Census tract boundary changes makes the decline in concentrated poverty even larger than in the Baseline measure, regardless of whether we use the Area-Weighted or the Target-

[5] We technically use the 2012-vintage, which includes a small set of Los Angeles County tract boundary changes made that year. The crosswalks and detailed descriptions of the methodology used to create them are available at `https://www.nhgis.org/geographic-crosswalks`.

Density-Weighted measure. Figure 3 plots all three measures of concentrated poverty for each city. The Baseline measure is displayed as a solid circle, the AW measure is a hollow circle, and the TDW measure is a hollow diamond. Cities are sorted from the most negative to the most positive Baseline measure. The leftward shift from the solid Baseline dots to the two reallocated measures indicates that both alternative measures consistently show a larger decline or, in the few cases of increasing concentrated poverty, a smaller increase. The AW and TDW measures are themselves nearly identical to each other, correlated at 0.99 across the 30 cities. The conclusion that the 2020 tract-boundary revision is not driving the measured decline in concentrated poverty is robust to the specific interpolation method chosen.

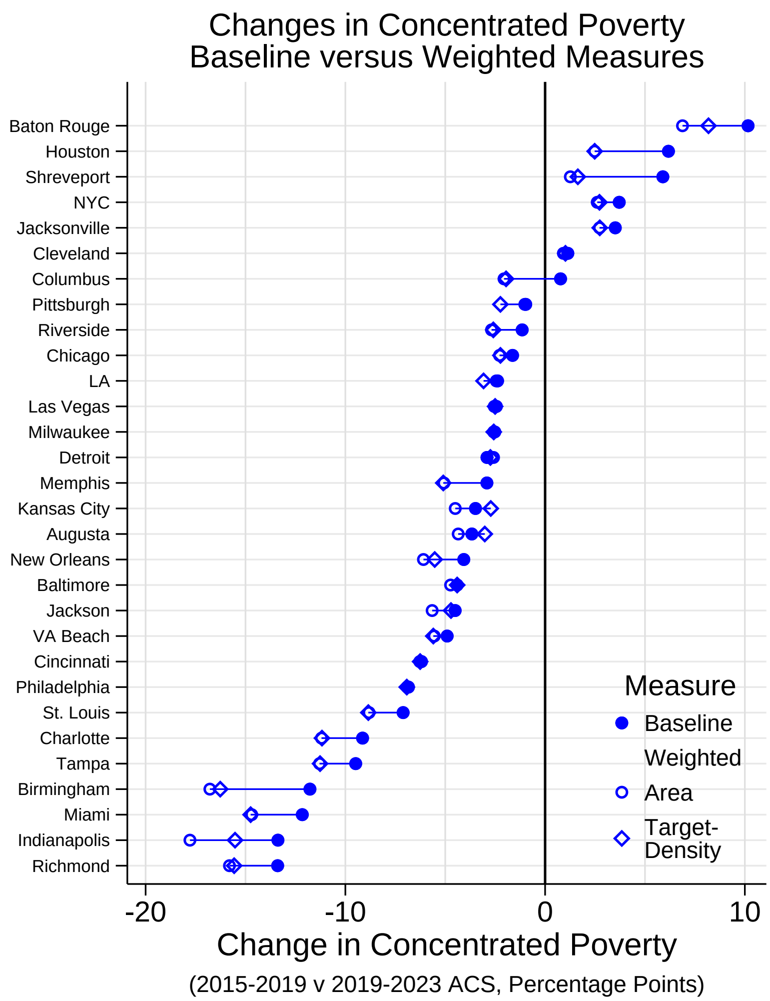


Figure 3: Baseline vs. Weighted Measures of the Change in Concentrated Poverty, by City

Note: Each row is one of the 30 cities in our sample, sorted from the most negative to the most positive Baseline measure of Concentrated Poverty. The solid blue circle represents the Baseline measure, the hollow blue circle represents the Area-Weighted measure, and the hollow blue diamond represents the Target-Density-Weighted measure.

## 3.2 The Stability of Labeling Conventions: Population Counts

Two secular trends in the poor populations of low-SES neighborhoods indicate uncertainty about the stability of racial labeling conventions pre- and post-2020.

A decline in the number of poor Black residents of low-SES neighborhoods began before the 2020 Census and continued its trend after the changes in the 2020 Census were implemented. The large red dots in Figure 4a show the levels of poor Black residents of Baltimore who were in Census tracts in the bottom quintile of neighborhood SES in each wave of the ACS. The red lines are best fit lines estimated separately via Ordinary Least Squares (OLS) for the period pre-2020 changes (2014-2017) and the period post-2020 changes (2018-2021). The number of poor Black residents in the lowest quintile of neighborhood SES declined from nearly 81 thousand in 2014 to just over 69 thousand in 2021. This decline continued at a similar pace in the four years of data before and after the 2020 changes in the measurement of race.

The data suggest a stable secular change in the poor Black population in low-SES neighborhoods. Appendix B shows that the pattern of stable change across the 2020 Census holds across our sample of segregated cities.

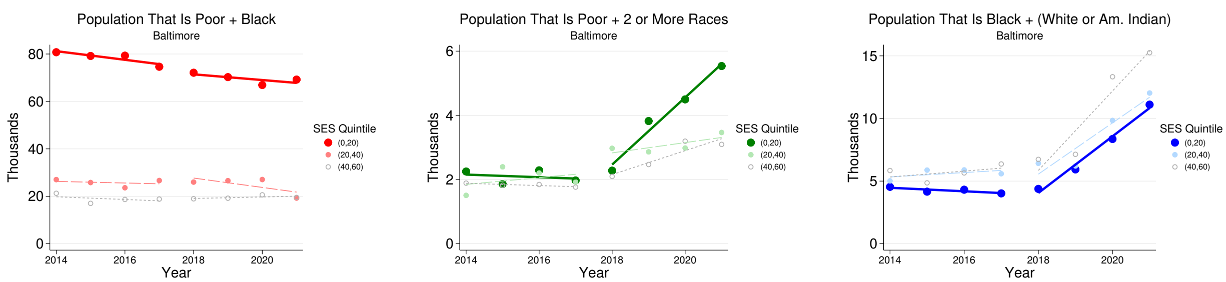


(a) Poor and Black

(b) Poor and Two or More Races

(c) Black and White or Black and American Indian

Figure 4: Population by Neighborhood SES Quintile and Race in Baltimore, by Vintage of the American Community Survey (ACS)

Note: These figures show population counts by race in the bottom three quintiles of neighborhood SES in Baltimore. These quintiles represent neighborhood SES levels of (0,20), (20,40), and (40,60).

However, secular changes in other relevant groups suggest that the changes in those groups' population are driven by changes in the way individuals are assigned racial labels. Beginning with the 2020 Decennial Census and American Community Survey (ACS), the Census Bureau modified its race and ethnicity questions and how it processed responses to those questions. There are several important changes, described in detail in Appendix A and Marks and Ríos-Vargas (2021). Broadly speaking, additional responses were allowed for write-in description of origins; the Bureau captured *200* characters from write-in responses in 2020 compared to *30* characters in prior censuses and the ACS; and the process for processing those write-in responses changed in substantive ways. As detailed in multiple analyses (Arias et al. (2025), Starr and Pao (2024), Reynolds (2021), National

Academies of Sciences, Engineering, and Medicine (2023)), we observe dramatic changes in the racial and ethnic composition of the US population that are driven by this new coding methodology. For example, between 2019 and 2020 the share of the overall population labeled as White alone declined from 72.0 percent to 62.7 percent and the Two or More Races population increased from 3.4 percent to 11.5 percent.

Here we show that, in contrast to the secular decline in poor Black residents, the secular trend in the number of poor residents in low-SES neighborhoods who are labeled as "Two or More Races" is not stable across the 2020 change in the measurement of race. This suggests the possibility that poor Black residents pre-2020 have experienced no change in neighborhood SES, but have simply been labeled with a different racial category post-2020.

Figure 4b shows the structural break in the "Two or More Races" group at the time of the 2020 change in the definition of race. The large green dots show the levels of poor minority residents of Baltimore in the "Two or More" racial group who were in Census tracts in the bottom quintile of neighborhood SES in each wave of the ACS. The green lines are best fit lines estimated separately via Ordinary Least Squares (OLS) for the period pre-2020 changes (2014-2017) and the period post-2020 changes (2018-2021). After the 2020 change in definition, starting at 2018 (or 2016-2020) in the figure, we see a sharp uptick in the number of poor "Two or More Races" residents of Baltimore's lowest SES neighborhoods. Appendix B shows that, just like the smooth trend for poor Black residents, this pattern of a structural break for poor "Two or More Race" residents holds across our sample of segregated cities.

The number of poor Black residents in the bottom quintile decline by 5.4 thousand between 2017 and 2021. The number of poor Two or More Races residents in the bottom quintile increased by 3.6 thousand between 2017 and 2021. Thus, the increase in the number of poor Two or More Races residents represents two thirds of the decrease in poor Black residents in low SES neighborhoods between 2017 and 2021.

If the secular trends in Figure 4 are driven by changes in labeling, one may wonder why the upward sloping lines in Figure 4b and 4c depict a smooth increase in the count of individuals in those specific racial categories. If the Census Bureau introduced a new processing methodology for race/ethnicity in 2020, we may have expected to observe a sharp discontinuity as opposed to a smooth increase.

The way the Bureau created the 5-year ACS estimates explains the lack of a discontinuity. The revised race/ethnicity questions and processing started with the responses received in 2020. Data collected in previous years was based on differently worded questions and processing methodology. Thus, when the Bureau produced the 2016-2020 ACS 5-year estimates, they combined four year's (2016-2019) of responses processed using the prior methodology with one year's (2020) responses using the new methodology. In each subsequent ACS release, additional years processed using the new methodology were added and years processed using the prior methodology were dropped. The 2020-2024 ACS released in December 2025 was the first 5-year ACS dataset with all race/ethnicity data processed using the new methodology (US Census Bureau (2020, 2023)).

### 3.3 The Stability of Labeling Conventions: Consequences for Measuring Concentrated Poverty

What are the consequences of uncertainty in the labeling of individuals for uncertainty in the measurement of concentrated poverty? We begin by formally defining our measure of segregation in a city, which we use interchangeably with concentrated poverty, in 2017,

$$\mathcal{S}_{2017} = \frac{\text{\# of Poor Black Residents with Neighborhood SES } \leq 10}{\text{\# of Poor Black Residents with Any Neighborhood SES}}. \tag{1}$$

Accurate measurement of changes in segregation depends on how individuals in the "Two or More Races" group Post-2020 would have been labeled Pre-2020. It is possible that the "Two or More Races" group in 2021 would have been labeled in a stable, identical way as "Two or More Races" Pre-2020. Under this assumption, we measure segregation in 2021 identically as we do in 2017, so that Equation 2 follows Equation 1. On the other hand, it is possible that the labeling of the "Two or More Races" group changed Pre- and Post-2020, where many of those would have been labeled as "Black" Pre-2020 are now labeled as "Two or More Races." Under this assumption, given an identical population in 2021 as in 2017, in order to measure segregation of the same individuals in 2021 as in 2017, we would need to measure segregation in 2021 as shown in Equation 3:

$$\mathcal{S}_{2021}^{\text{Stable}} = \frac{\text{\# of Poor Black w SES } \leq 10}{\text{\# of Poor Black w Any SES}} \tag{2}$$

$$\mathcal{S}_{2021}^{\text{Changing}} = \frac{\text{\# of Poor Black w SES } \leq 10 \;+\; \text{2017-2021 change in \# of Poor 2 or More Races w SES } \leq 10}{\text{\# of Poor Black w Any SES} \;+\; \text{2017-2021 change in \# of Poor 2 or More Races w Any SES}} \tag{3}$$

We define measures of changes in segregation in terms of their assumption about the labeling of the "Two or More Races" group Pre- and Post-2020:

$$\Delta^{\text{Stable}} = \mathcal{S}_{2021}^{\text{Stable}} - \mathcal{S}_{2017}$$

$$\Delta^{\text{Changing}} = \mathcal{S}_{2021}^{\text{Changing}} - \mathcal{S}_{2017}$$

For our ensuing empirical work we calculate the 2017-2021 change in the number of poor "Two or More Races" individuals using regressions of the poor "Two or More Races" population on year between 2017 and 2021 separately in neighborhoods above and below the 10th percentile of SES and then multiplying those regression coefficients by four.

Figure 5 uses these alternate assumptions about labeling conventions to show the range of uncertainty when measuring changes in concentrated poverty. The figure orders cities by the magnitude of their difference under the two labeling conventions. We see three broad patterns.

In the most affected cities, shown in navy blue, different labeling conventions would completely alter the measured changes in concentrated poverty. These cities include Houston, New York City (NYC), Riverside, Los Angeles (LA), Las Vegas, Pittsburgh, Chicago, Milwaukee, Kansas City, Tampa Bay, and Miami. As an example, assuming a stable labeling convention implies that there was a large decline in concentrated poverty in Milwaukee, while assuming a changing labeling convention implies there was a negligible increase in concentrated poverty there.

Another set of affected cities would experience a non-trivial, quantitative change to the interpretation of their trend in concentrated poverty. This group of cities, shown in blue, comprises Cleveland, Columbus, New Orleans, Baltimore, Virginia Beach, Charlotte, and Indianapolis. For example, assuming a stable labeling convention implies that there was a large decline in concentrated poverty in New Olreans, while assuming a changing labeling convention implies that the magnitude of the decline was much smaller.

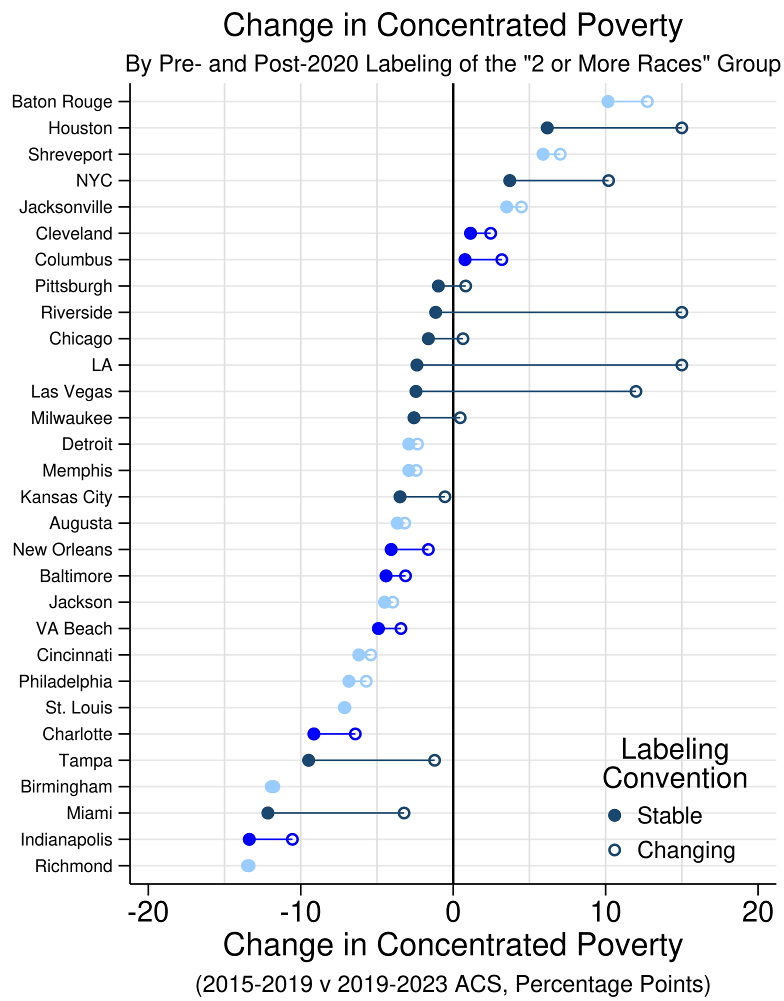


Figure 5: Change in Concentrated Poverty by Labeling Convention, by City

Note: This figure shows changes in concentrated poverty, or the percent of poor Black residents of a city living in a neighborhood with socioeconomic status (SES) less than the 10th percentile of neighborhoods in the US. The blue dots show changes by city between the 2015-2019 and 2019-2023 waves of the American Community Survey (ACS). The solid blue dots assume a stable labeling convention, that individuals labeled "Black" pre-2020 would be identically labeled as "Black" post-2020. The hollow blue dots assume a changing labeling convention, in which increases in the "Two or More Races" population are entirely due to individuals labeled "Black" pre-2020 having their labels switched to "Two or More Races" post-2020. City-level changes in poor "Two or More Races" population are estimated via OLS regressions on the 2015-2019 through 2019-2023 waves of the ACS.

Finally, there is a set of cities shown in light blue for which the interpretation is not altered by the labeling convention. This group of cities comprises Baton Rouge, Shreveport, Jacksonville, Detroit, Memphis, Augusta, Jackson, Cincinnati, Philadelphia, St. Louis, Birmingham, and Richmond. St. Louis is an example where both labeling conventions indicate the city experienced a massive decline in concentrated poverty.

We note that Hispanic migration is likely to play a role in many of the most-affected cities. As a result, we consider a robustness check in which only a fraction of the increase in "Two or More Races" individuals were due to labeling changes of "Black" residents (in our exercise, that fraction was 1/5). Even in this scenario, a completely altered qualitative picture would still be painted for Los Angeles, Riverside, Houston, and Las Vegas.

# 4 Conclusion

Given the history of race in the US, there is a normative question of how race should be defined today (O'Flaherty (2016), Darity and Lefebvre (2024)). This note studied the related positive (empirical) question of how to accurately characterize the US under different measures of race and ethnicity, especially those used before and after the changes made by the US Census Bureau in 2020 (Arias et al. (2025)).

We showed that changes to the measurement of race beginning in 2020 generate a source of conceptual uncertainty when interpreting official statistics (Manski (2015)). The uncertainty we document makes it difficult to accurately characterize neighborhood inequality. One ultimate goal of social scientists is to characterize neighborhoods in a way that captures their causal effects on residents (Aliprantis et al. (2024), Chetty et al. (2020), Noelke et al. (2020)). The more preliminary goal of accurately characterizing the current residents of neighborhoods is itself non-trivial. For example, measures of income segregation mechanically respond to changes in the population distribution of income, complicating the interpretation of such measures (Watson (2009)). Similarly, sampling variation obfuscates the measurement of residential segregation by either race (Napierala and Denton (2017)) or income (Logan et al. (2018), Reardon et al. (2018)) as well as the rankings of neighborhoods (Mogstad et al. (2024)). Our results add another source of uncertainty to even this preliminary goal.

We see two key paths for future research into understanding the uncertainty we document in this note. First, we note that the third (blue) column in Appendix B shows that for many metros, the counts of population that are Black in combination with either white or American Indian jumped discretely between the 2017-2021 and 2018-2022 waves of ACS, which were the second and third waves of the ACS for which the 2020 changes were included. Is this pattern due to obstacles COVID-19 created for the Census Bureau and the ACS? Relatedly, future research could disentangle the underlying reasons that the racial labels attached to individuals changed. It is possible that there were genuine changes in respondents' multiracial identification that coincided in timing with the measurement changes made by the Census in 2020. However, the evidence indicates that the impact of such identification changes is minor relative to the Census' changes (Ventura and Flores (2025), Starr and Pao (2024)). Focusing therefore on the changes in the Census, there are many purposes for which it will be useful to understand how much of the differences in racial labeling stems from the introduction of the new race question itself versus changes to editing and recoding rules.

Optimal policy and societal outcomes hinge critically upon whether the massive recent decline in concentrated poverty in US cities is due to changes in the neighborhoods where poor Black residents live or else is simply due to changes in the way racial labels are assigned to individuals. This note showed that, unfortunately, in more than half of the cities in our sample we are not able to confidently interpret the measured decline in concentrated poverty.

Overall poverty declined considerably over the time period studied in this note (Appendix Figure 11). Increases in income could break the concentration of poverty by allowing poor households to move out of high-poverty neighborhoods (Garin et al. (2025)). If the measured decline in concentrated poverty reflects this pattern, then the most effective strategy for combating the effects of concentrated poverty may simply be overall economic growth. This could be seen as consistent with the trends in concentrated poverty over recent decades (Jargowsky (2003), Dwyer (2012), Kneebone and Nadeau (2016)).

On the other hand, we know that racial inequality is highly persistent in the US. There has been no change in the ratios of the mean Black and white income or wealth since the 1960s (Aliprantis et al. (2026), Derenoncourt et al. (2024)) and Black men's contemporaneous earnings are in the same location in the distribution as in 1940 (Bayer and Charles (2018)). This note has demonstrated that in many cities the large decline in overall poverty since 2014 could have had a limited impact on the neighborhood conditions experienced by the poor Black residents of high poverty neighborhoods.

More broadly, we emphasized that the problem documented in this note is not specific to the study of concentrated poverty or to data from the Census or ACS. Any study of trends in racial inequality must grapple with the implications of changes over time to the definition of race itself. The 2020 changes made by the US Census Bureau are a salient and timely example of this issue.

## References


Aliprantis, D., D. R. Carroll, and E. R. Young (2026). The dynamics of the racial wealth gap. *International Economic Review*. Forthcoming.

Aliprantis, D. and S. DeLuca (2025). Unlocking opportunity: The remarkable success of the Baltimore Regional Housing Partnership. *Mimeo., ENS de Lyon*.

Aliprantis, D., H. Martin, and K. Tauber (2024). What determines the success of housing mobility programs? *Journal of Housing Economics 65*, 102009.

Arias, E., C. A. Liebler, M. A. Garcia, and R. Sáenz (2025). Data impacts of changes in US Census Bureau procedures for race and ethnicity data. *SSM - Population Health 29*, 101742.

Bayer, P. and K. K. Charles (2018). Divergent paths: A new perspective on earnings differences between black and white men since 1940. *The Quarterly Journal of Economics 133*(3), 1459–1501.

Chetty, R., J. N. Friedman, N. Hendren, M. R. Jones, and S. R. Porter (2020). The Opportunity Atlas: Mapping the childhood roots of social mobility. *Mimeo., Opportunity Insights*.

Darity, Jr., W. A. and S. Lefebvre (2024). Data collection without definitions. In R. Akee, L. F. Katz, and M. Loewenstein (Eds.), *Race, Ethnicity, and Economic Statistics for the 21st Century*, Chapter 2. University of Chicago Press.

Darrah, J. and S. DeLuca (2014). "Living here has changed my whole perspective": How escaping inner-city poverty shapes neighborhood and housing choice. *Journal of Policy Analysis and Management 33*(2), 350–384.

Davis, D. B. (2006). *Inhuman Bondage: The Rise and Fall of Slavery in the New World*. Oxford University Press.

DeLuca, S. and P. Rosenblatt (2017). Walking away from *The Wire*: Housing mobility and neighborhood opportunity in Baltimore. *Housing Policy Debate 27*(4), 519–546.

Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2024). Wealth of two nations: The US racial wealth gap, 1860–2020. *The Quarterly Journal of Economics 139*(2), 693–750.

Duncan, B. and S. J. Trejo (2011). Intermarriage and the intergenerational transmission of ethnic identity and human capital for Mexican Americans. *Journal of Labor Economics 29*(2), 195–227.

Dwyer, R. E. (2012). Contained dispersal: The deconcentration of poverty in US metropolitan areas in the 1990s. *City & Community 11*(3), 309–331.

Garin, A., E. Jenkins, E. Mast, and B. Stuart (2025). Dynamic individuals, static neighborhoods: Migration, earnings changes, and concentrated poverty. *Mimeo., University of Notre Dame*.

Glaeser, E. L., J. Gyourko, and B. Neiszner (2025). Measuring neighborhood change: The issue of ex post borders. *NBER Working Paper 34238*.

Jargowsky, P. (2003). Stunning progress, hidden problems: The dramatic decline of concentrated poverty in the 1990s. In A. Berube, B. Katz, and R. E. Lang (Eds.), *Redefining Urban and Suburban America: Evidence from Census 2000*, Chapter 5. Washington, DC: The Brookings Institution.

Jones, N., E. Jensen, K. Battle, and R. Marks (2024). Measuring the racial and ethnic composition and diversity of the United States population: Historical challenges and contemporary opportunities. In R. Akee, L. F. Katz, and M. Loewenstein (Eds.), *Race, Ethnicity, and Economic Statistics for the 21st Century*, Chapter 1. University of Chicago Press.

Jones, N., R. Marks, R. Ramirez, and M. Ríos-Vargas (2021, August 12). *2020 Census Illuminates Racial and Ethnic Composition of the Country*. America Counts: Stories Behind the Numbers: US Census Bureau. Retrieved from https://www.census.gov/library/stories/2021/08/improved-race-ethnicity-measures-reveal-united-states-population-much-more-multiracial.html.

Kneebone, E. and C. A. Nadeau (2016). The resurgence of concentrated poverty in America: Metropolitan trends in the 2000s. In K. B. Anacker (Ed.), *The New American Suburb: Poverty, Race, and the Economic Crisis*. Routledge.

Light, A. and A. Nandi (2007). Identifying race and ethnicity in the 1979 National Longitudinal Survey of Youth. *Population Research and Policy Review 26*(2), 125–144.

Logan, J. R., A. Foster, J. Ke, and F. Li (2018). The uptick in income segregation: Real trend or random sampling variation? *American Journal of Sociology 124*(1), 185–222.

Manski, C. F. (2015). Communicating uncertainty in official economic statistics: An appraisal fifty years after Morgenstern. *Journal of Economic Literature 53*(3), 631–653.

Manson, S., J. Schroeder, D. V. Riper, K. Knowles, T. Kugler, F. Roberts, , and S. Ruggles (2024). *IPUMS National Historical Geographic Information System: Version 19.0 [dataset]*. Minneapolis, MN: IPUMS.

Marks, R. and M. Ríos-Vargas (2021, August 3). *Improvements to the 2020 census race and hispanic origin question designs, data processing, and coding procedures*. Random Samplings Blog: US Census Bureau. Retrieved from https://www.census.gov/newsroom/blogs/random-samplings/2021/08/improvements-to-2020-census-race-hispanic-origin-question-designs.html.

Mason, P. L. (2023). *The Economics of Structural Racism: Stratification Economics and US Labor Markets*. Cambridge University Press.

Mogstad, M., J. P. Romano, A. M. Shaikh, and D. Wilhelm (2024). Inference for ranks with applications to mobility across neighbourhoods and academic achievement across countries. *Review of Economic Studies 91*(1), 476–518.

Napierala, J. and N. Denton (2017, 01). Measuring residential segregation with the ACS: How the margin of error affects the dissimilarity index. *Demography 54*(1), 285–309.

National Academies of Sciences, Engineering, and Medicine (2023). *Assessing the 2020 Census: Final Report*. Washington, DC: The National Academies Press.

Noelke, C., N. McArdle, M. Baek, N. Huntington, R. Huber, E. Hardy, and D. Acevedo-Garcia (2020). *Child Opportunity Index 2.0 Technical Documentation.*

O'Flaherty, B. (2016). *Race May be Pseudo-Science, But Economists Ignore it at their Peril* (INET Conference on the Economics of Race ed.). Detroit, MI: Institute for New Economic Thinking.

OMB (1997, October 30). Revisions to the standards for the classification of federal data on race and ethnicity. *Federal Register 62*(210), 58782–58790.

Reardon, S. F., K. Bischoff, A. Owens, and J. B. Townsend (2018). Has income segregation really increased? Bias and bias correction in sample-based segregation estimates. *Demography 55*(6), 2129–2160.

Reynolds, L. (2021). Measuring race and ethnicity with the 2020 census redistricting data. *Mimeo., Cornell University Program on Applied Demographics*.

Ruggles, S., S. Flood, M. Sobek, D. Backman, G. Cooper, J. A. R. Drew, S. Richards, R. Rodgers, J. Schroeder, and K. C. Williams (2025). *IPUMS USA: Version 16.0 [dataset].* Minneapolis, MN: IPUMS.

Saperstein, A. and A. M. Penner (2012). Racial fluidity and inequality in the United States. *American Journal of Sociology 118*(3), 676–727.

Schroeder, J., D. V. Riper, S. Manson, G. Cooper, Z. Krause, T. Kugler, T. Zhu, and S. Ruggles (2026). *IPUMS National Historical Geographic Information System: Version 21.0 [dataset].* Minneapolis, MN: IPUMS.

Schroeder, J. P. (2007). Target-density weighting interpolation and uncertainty evaluation for temporal analysis of census data. *Geographical Analysis 39*(3), 311–335.

Shrider, E. A. and C. Bijou (2025, September 9). *Poverty in the United States: 2024*. Washington, DC: US Census Bureau. Report Number: P60-287.

Starr, P. and C. Pao (2024). The multiracial complication: The 2020 census and the fictitious multiracial boom. *Sociological Science 11*(40), 1107–1123.

US Census Bureau (2020). *American Community Survey and Puerto Rico Community Survey 2020 Subject Definitions.* Retrieved from https://assets.nhgis.org/original-data/acs/2020_ACSSubjectDefinitions.pdf.

US Census Bureau (2023). *American Community Survey and Puerto Rico Community Survey 2023 Subject Definitions.* Retrieved from https://assets.nhgis.org/original-data/acs/2023_ACSSubjectDefinitions.pdf.

US Census Bureau (2024, December 20). *Research to Improve Data on Race and Ethnicity.* Retrieved from https://www.census.gov/about/our-research/race-ethnicity.html.

Ventura, I. M. and R. D. Flores (2025). The "rise" of multiracials? Examining the growth in multiracial identification in the 2020 us census. *RSF: The Russell Sage Foundation Journal of the Social Sciences 11*(1), 44–64.

Washington, H. A. (2006). *Medical Apartheid: The Dark History of Medical Experimentation on Black Americans from Colonial Times to the Present.* Harlem Moon Broadway.

Watson, T. (2009). Inequality and the measurement of residential segregation by income in American neighborhoods. *Review of Income and Wealth 55*(3), 820–844.

Williams-Baron, E. and A. Saperstein (2026). A roadmap for inequality research: Transparency, Intersectionality, and Multiple Measures of Race. *Sociological Science 13*, 825–863.

Wilson, W. J. (1987). *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy.* University of Chicago press.

Zuberi, T. (2001). *Thicker Than Blood: How Racial Statistics Lie.* University of Minnesota Press.

# A Appendix: The Census Changed How It Processed Race Responses in 2020

Race and ethnicity classifications reported in official U.S. federal statistics have changed over time, with the most recent change introduced in 1997. The Office of Management and Budget updated Statistical Policy Directory Number 15 to add Native Hawaiian or Other Pacific Islander (NHOPI) as its own category and allowed individuals to select two or more race categories if they felt that matched their self-identification. The introduction of two more categories complicated temporal analysis of statistical data. To help researchers bridge this change in categories, federal agencies and researchers provided guidance in several forms.

Beginning with the 2020 Decennial Census and American Community Survey (ACS), the Census Bureau modified its race and ethnicity questions and how it processed responses to those questions.[6] For the race and ethnicity questions, the Bureau provided space for individuals identifying as any race (White, Black or African American, American Indian or Alaska Native, Asian, Native Hawaiian or Other Pacific Islander, or Some Other Race) to write in their origins (e.g., German, Irish, Ethiopian, Blackfeet Tribe, Hmong, Tongan, etc.). In the 2000 and 2010 Decennial Censuses and all ACS surveys through 2019, no write-in spaces were provided for the White or Black or African American categories; instead, the form just provided checkboxes. Providing write-in spaces for Whites and Black or African Americans results in a tenfold increase in the number of write-ins from 38 million in 2010 to 336 million in 2020 (National Academies of Sciences, Engineering, and Medicine (2023)).

Write-in responses must be processed by the Census Bureau in order to assign the racial categories to each person. Four important changes to the Bureau's processing methodology resulted in significant changes in the racial composition of the US. First, the Bureau captured *200* characters from write-in responses in 2020 compared to *30* characters in prior censuses and the ACS. Next, the Bureau considered up to six responses in 2020 compared with up to two responses in prior censuses. Then, the Bureau used a single code list for Hispanic ethnicities and race categories in 2020 compared with separate lists for each write-in response in prior censuses. Finally, the Bureau codes responses from left to right. In prior censuses, if a respondent wrote in a detailed Hispanic origin and two detailed racial origins into the race question, the Bureau privileged the detailed racial origin responses when coding the data. In 2020, the Bureau could have coded all three responses.

A simplified example of this coding change helps illustrate its impact on the data. Suppose someone wrote in the following string to both the 2010 and 2020 Census race questions – "Cuban, Thai, Filipino". In 2010, this person would be assigned to the Asian racial category. The Bureau privileged the two detailed Asian race responses over the Hispanic ethnicity group "Cuban". In 2020, this person would be assigned to both the Some Other Race and the Asian racial categories. Thus, they would be considered multiracial in 2020.

As detailed in multiple analyses (Arias et al. (2025), Ventura and Flores (2025), Starr and

[6] Information in the three following paragraphs is largely drawn from Marks and Ríos-Vargas (2021).

Pao (2024), Reynolds (2021), National Academies of Sciences, Engineering, and Medicine (2023)), we observe dramatic changes in the racial and ethnic composition of the US population that are partially (or mostly) driven by this new coding methodology. Between 2010 and 2020 the percentage of the White alone population declined from 72.4 percent to 61.6 percent and the percentage of the Two or More Race population increased from 2.9 percent to 10.2 percent. While we expect changes in the racial composition over a 10-year time period, we observe dramatic changes in the 2019 and 2020 ACS data. Between 2019 and 2020 the share of the overall population labeled as White alone declined from 72.0 percent to 62.7 percent and the Two or More Races population increased from 3.4 percent to 11.5 percent.

Changes to the Census Bureau's race and ethnicity coding provides a more nuanced and accurate view of the racial composition of the United States going forward (Jones et al. (2021)). But, these changes also disrupt researchers' ability to analyze change over time. We can no longer disentangle actual demographic change from changes introduced by the new race and ethnicity coding system.

# B Appendix: Secular Trends by Race

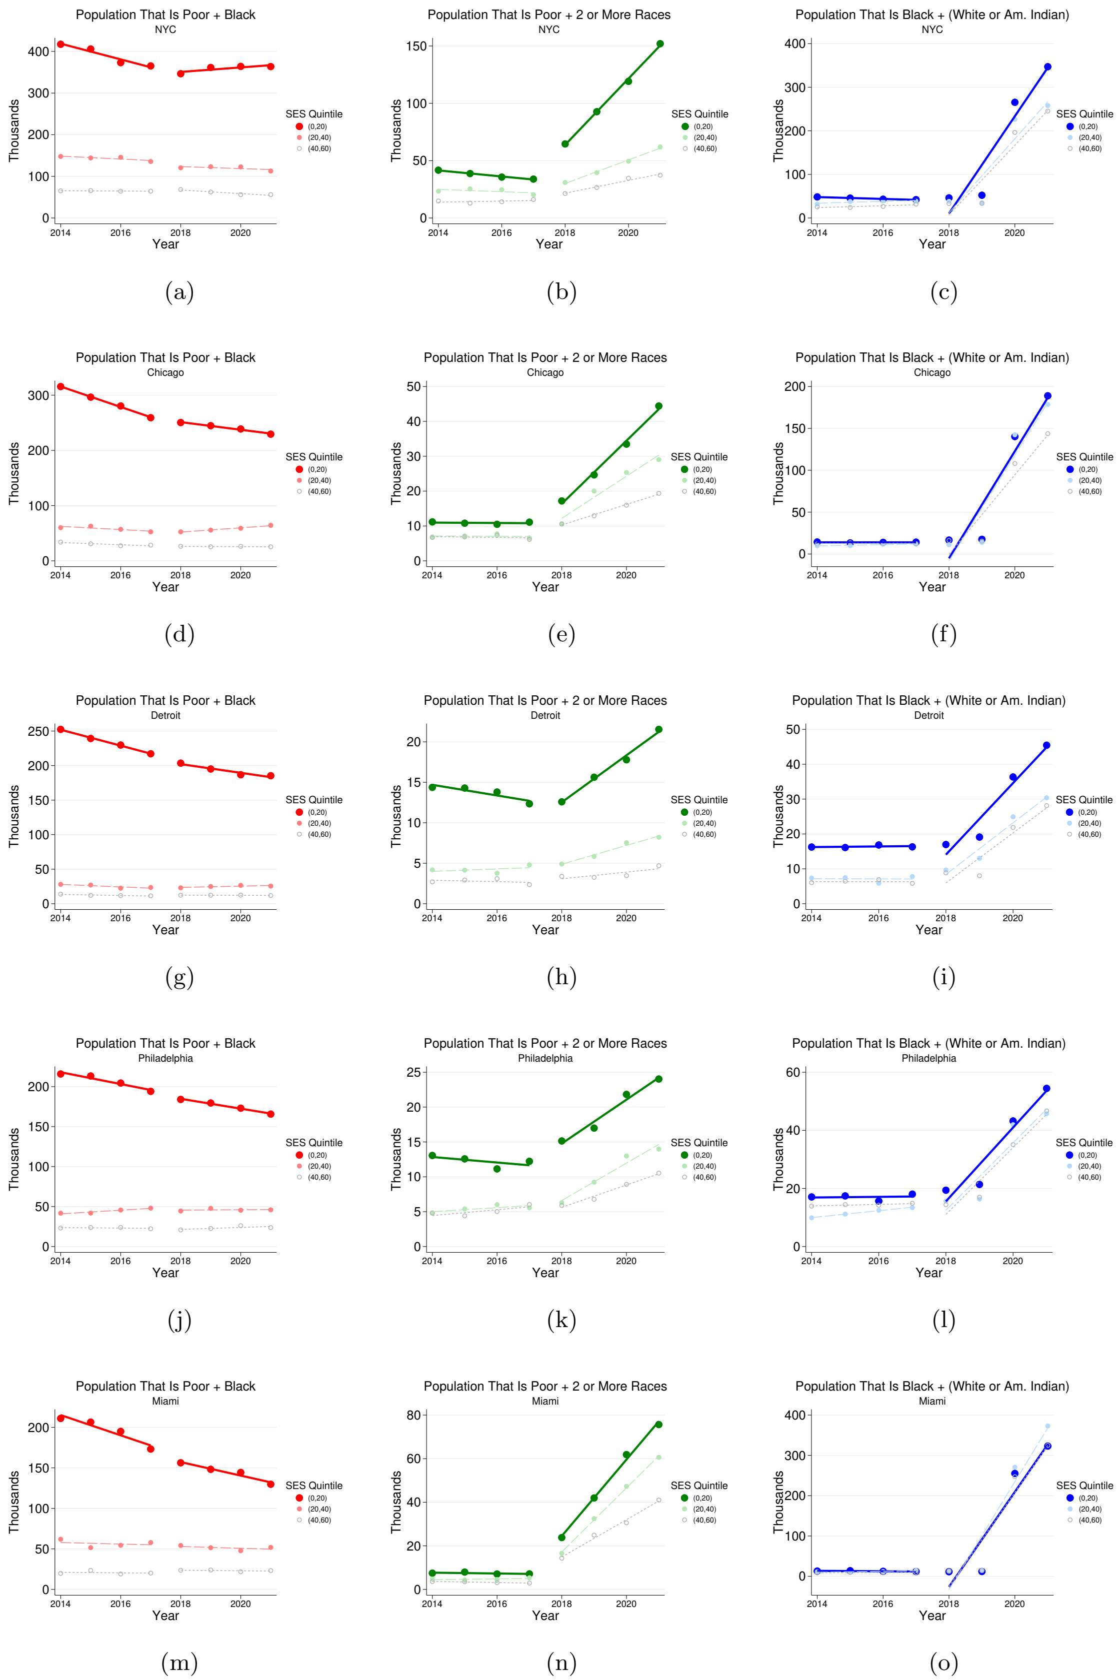




Figure 1: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

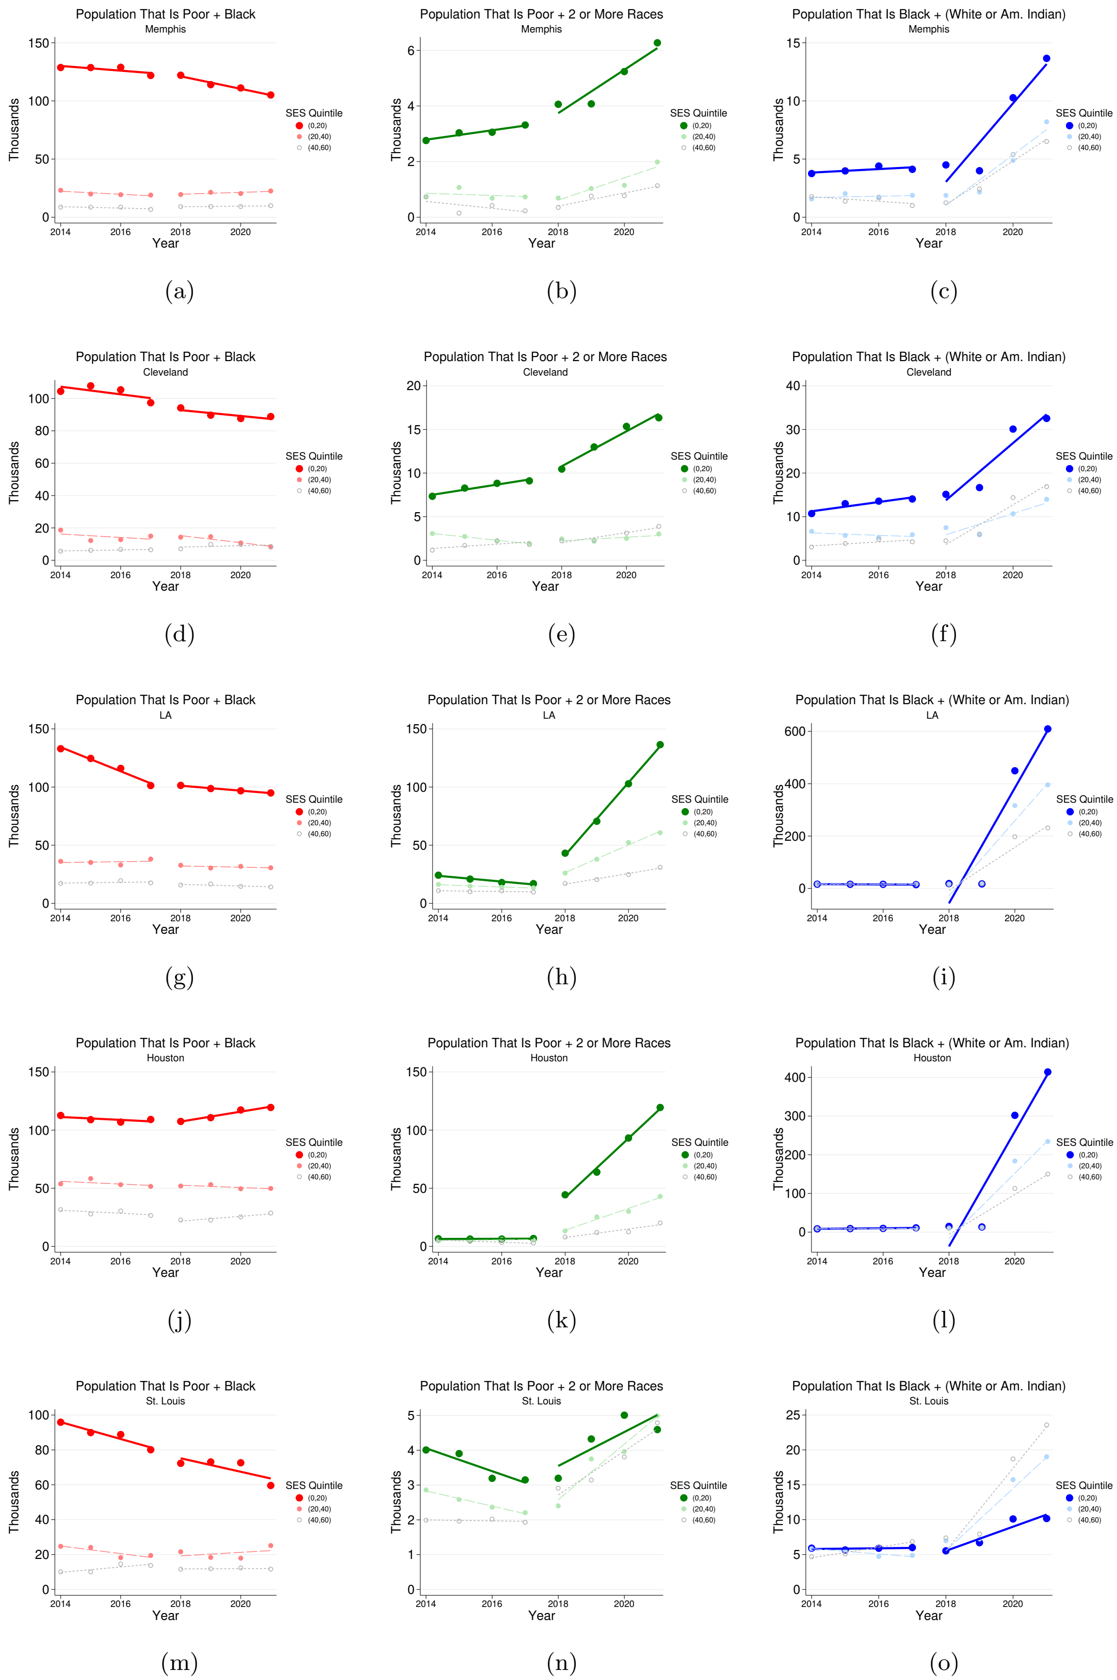


Figure 2: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

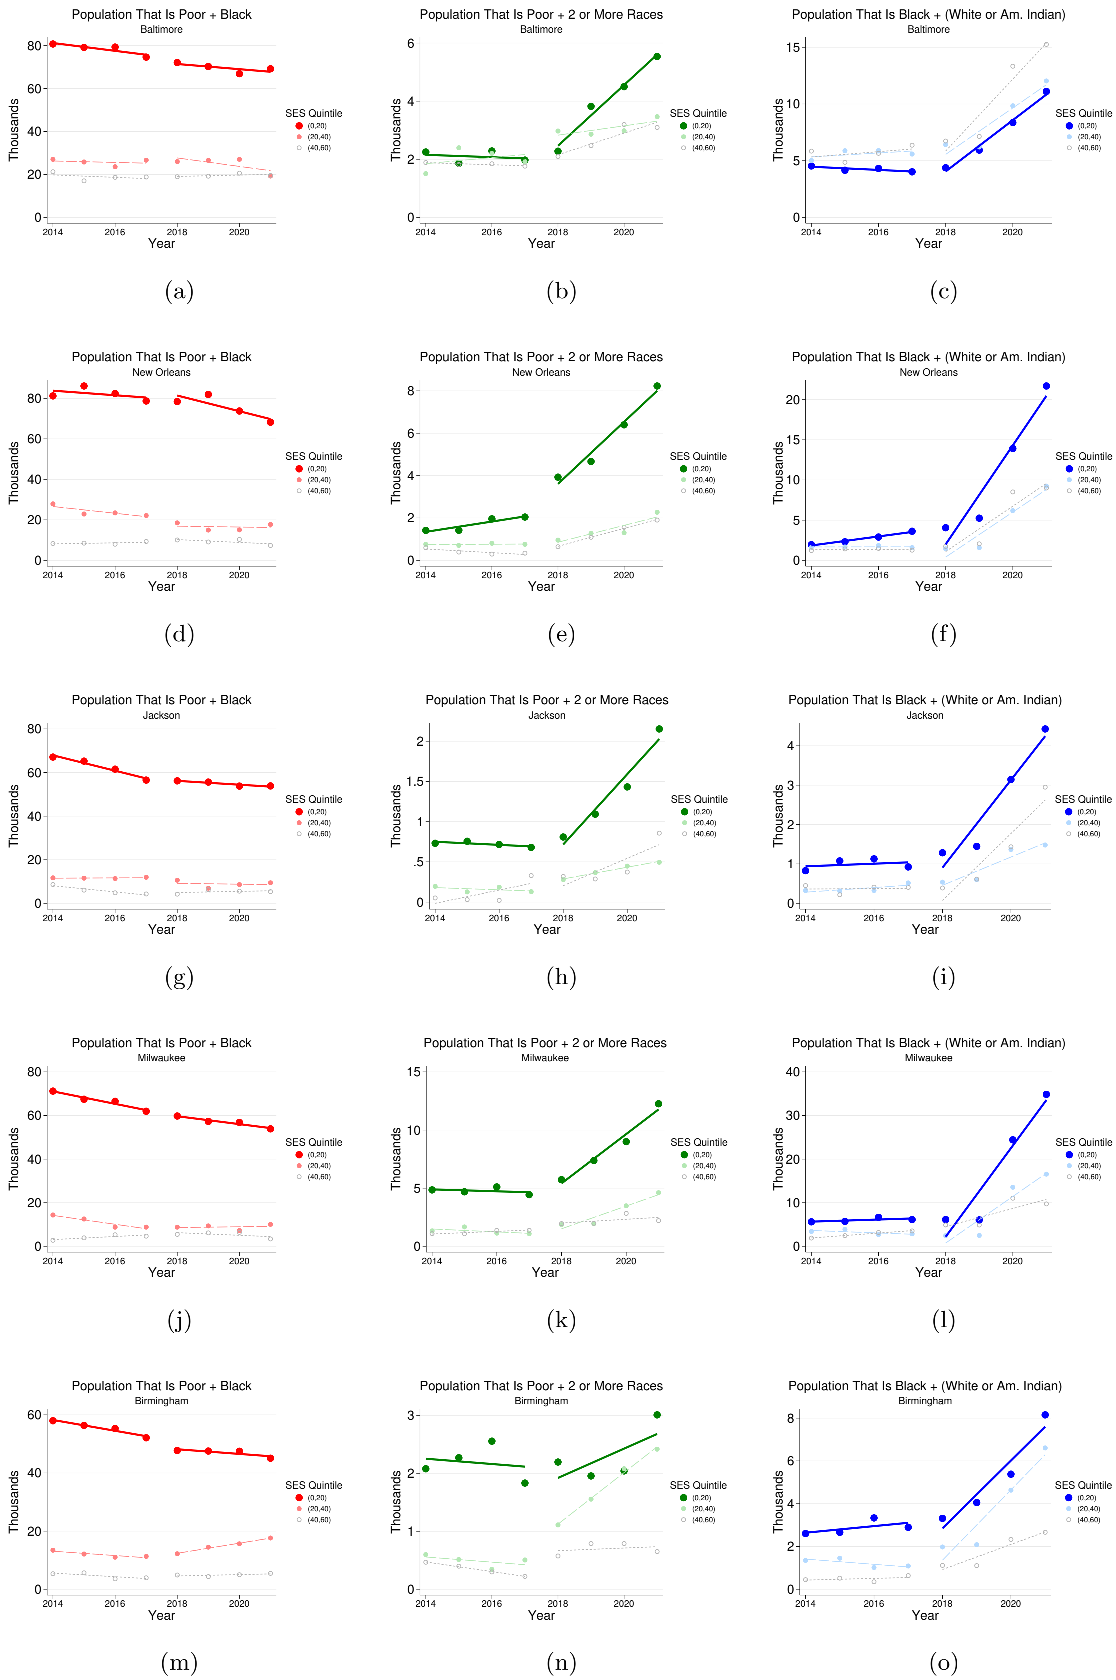


Figure 3: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

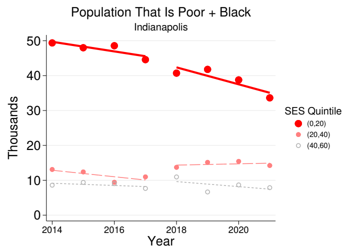


(a)

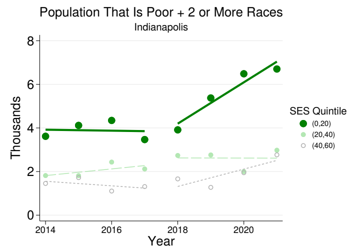


(b)

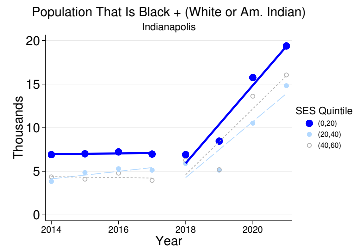


(c)

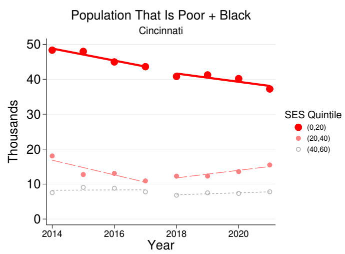


(d)

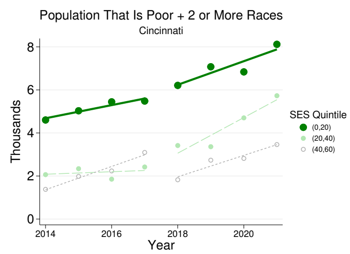


(e)

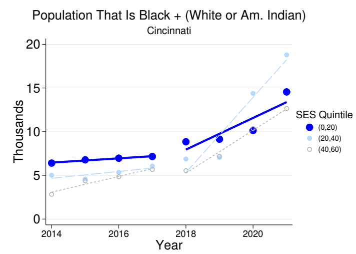


(f)

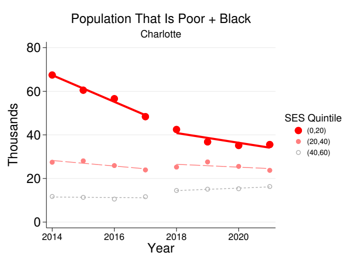


(g)

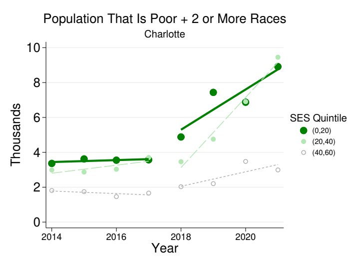


(h)

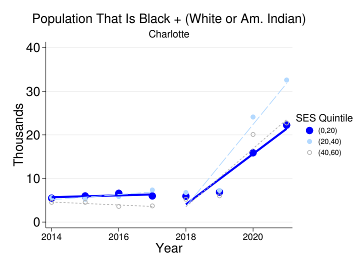


(i)

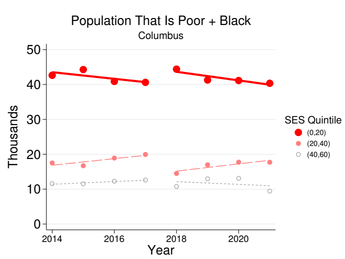


(j)

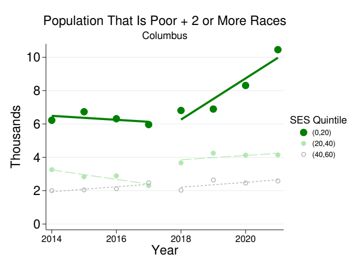


(k)

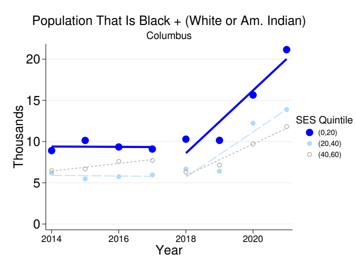


(l)

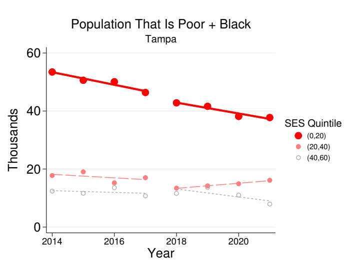


(m)

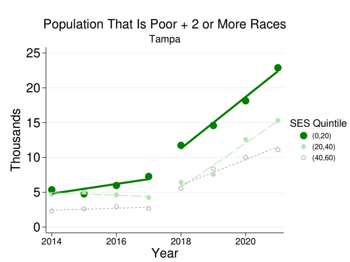


(n)

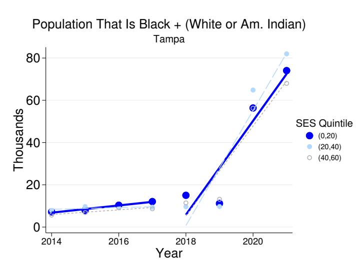


(o)

Figure 4: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

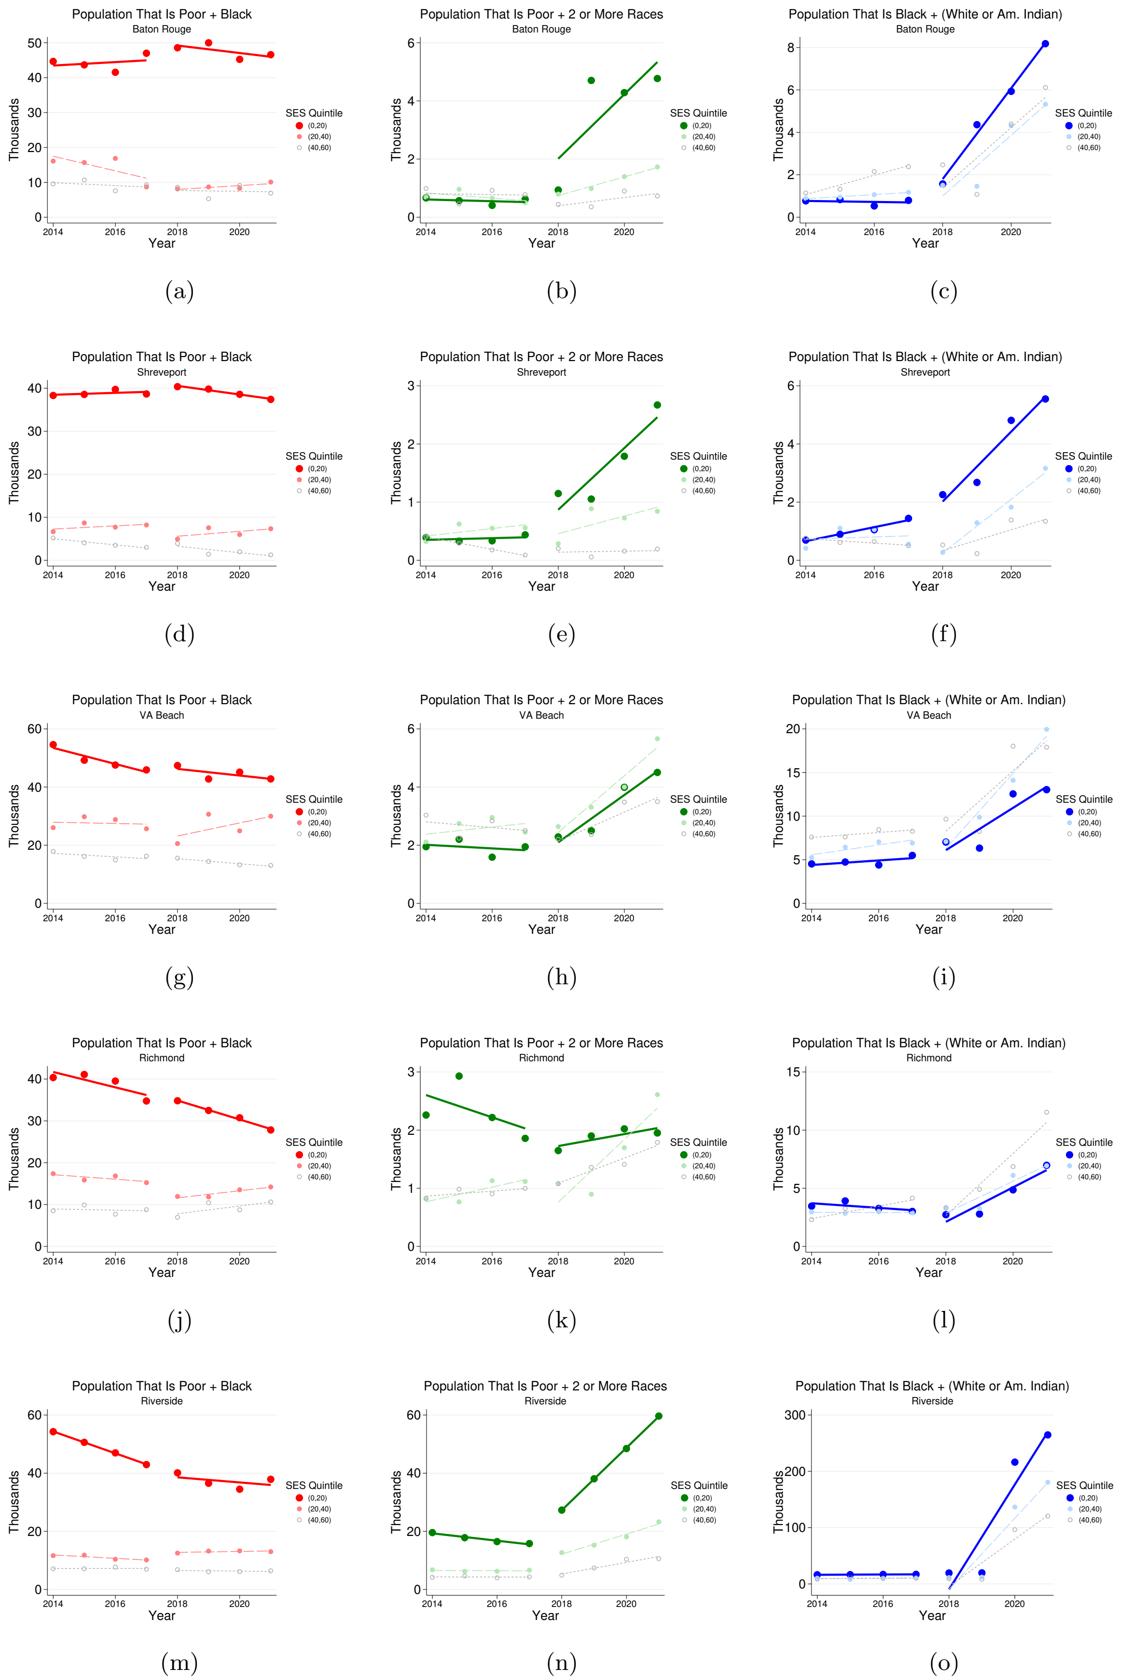


Figure 5: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

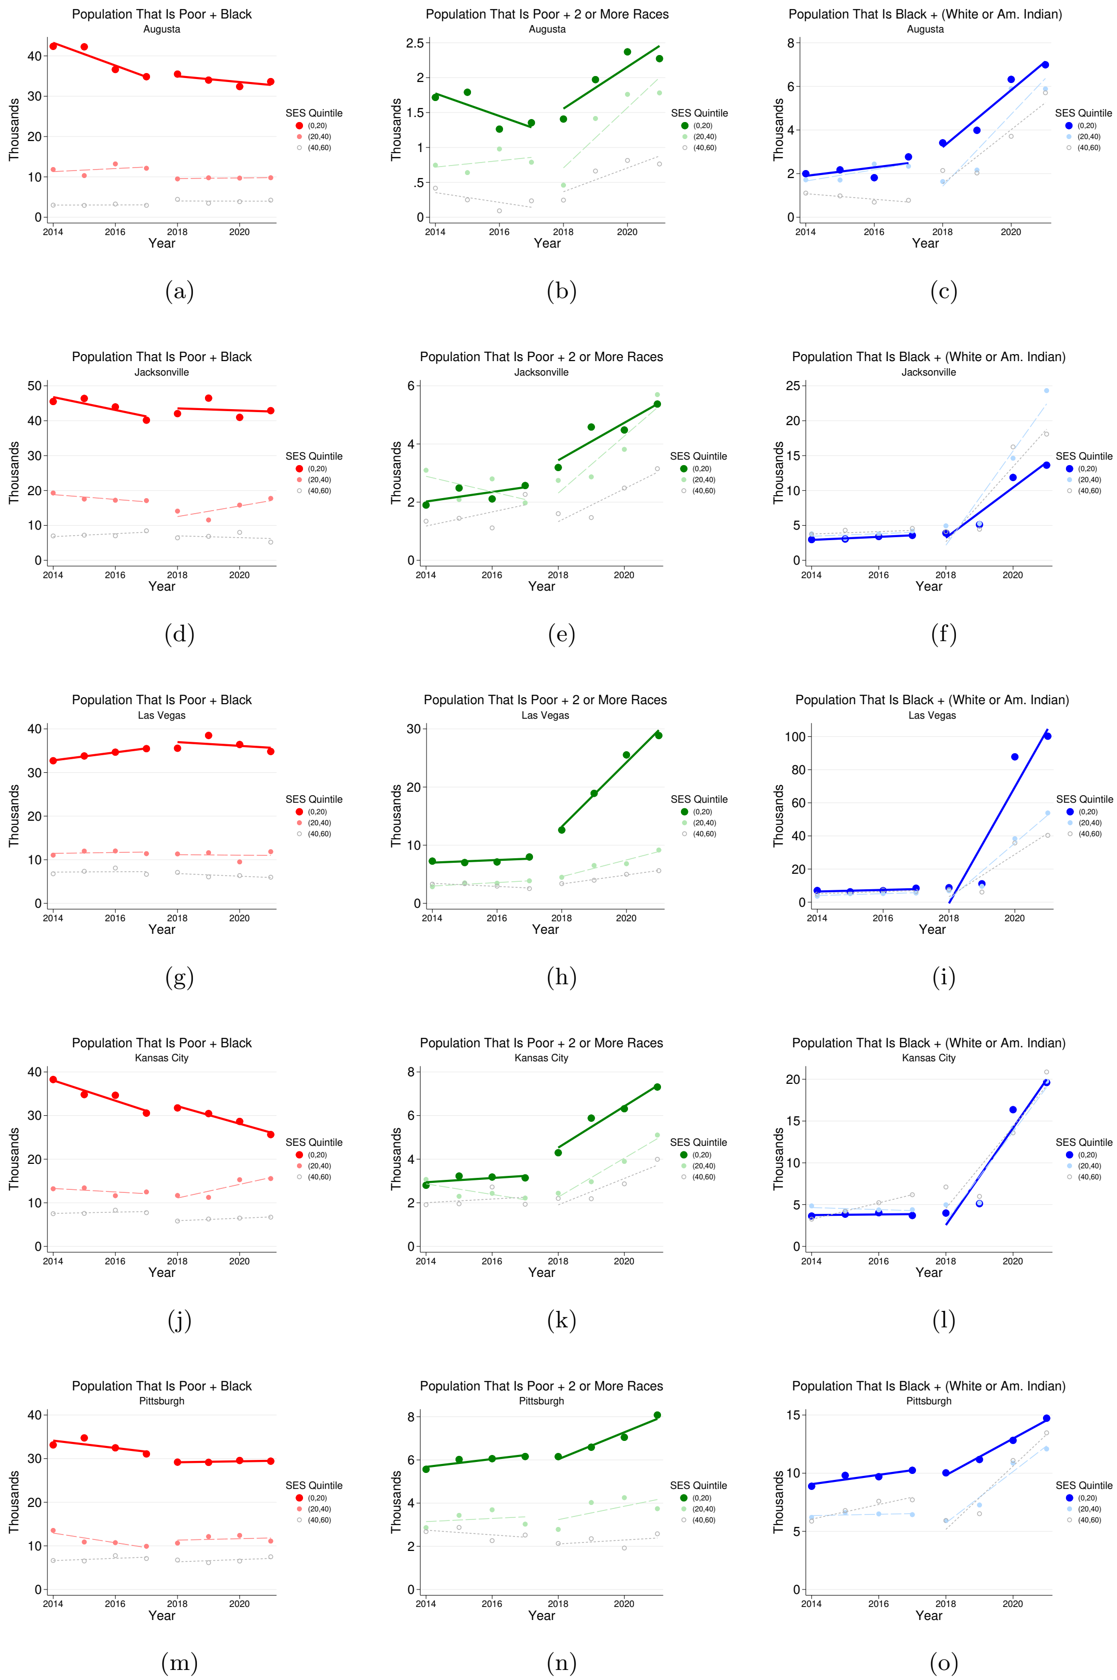


Figure 6: Population by Neighborhood SES Quintile and Race, by City and Vintage of the American Community Survey (ACS)

# C Appendix: Racial Shares in the Bottom Quintile of Neighborhood SES

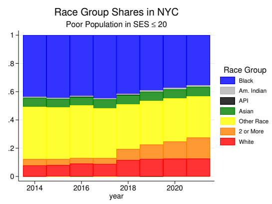


(a)

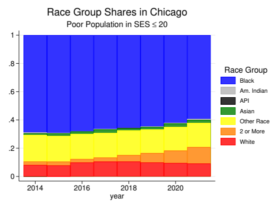


(b)

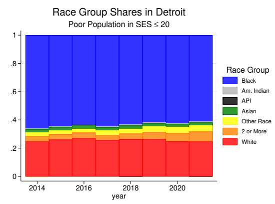


(c)

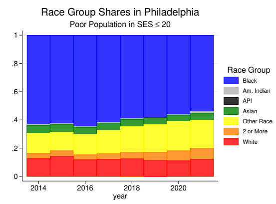


(d)

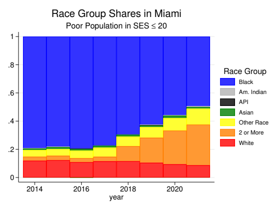


(e)

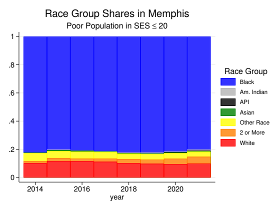


(f)

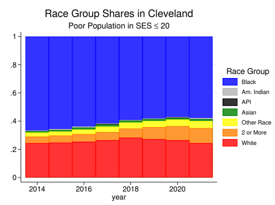


(g)

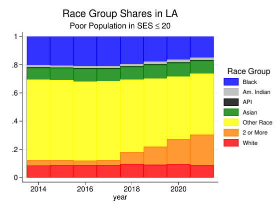


(h)

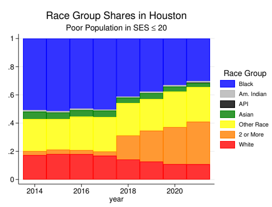


(i)

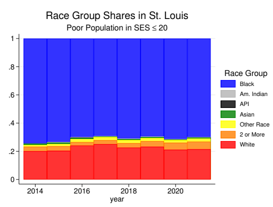


(j)

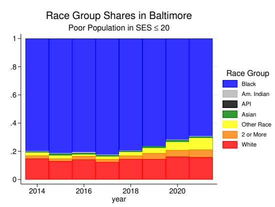


(k)

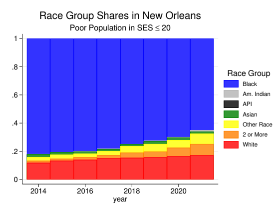


(l)

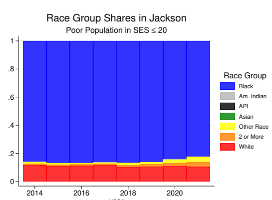


(m)

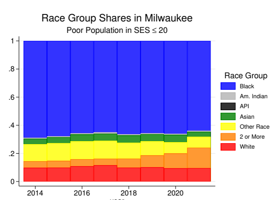


(n)

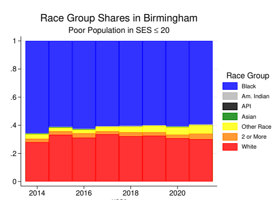


(o)

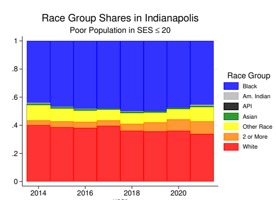


(p)

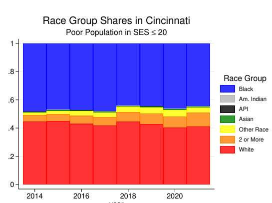


(q)

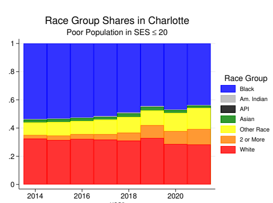


(r)

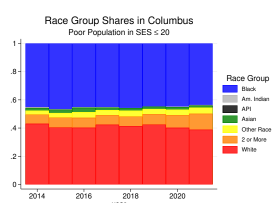


(s)

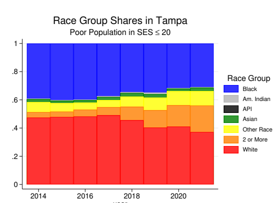


(t)

Figure 7: Racial Shares in the Bottom Quintile of Neighborhood SES, by City and Vintage of the American Community Survey (ACS)

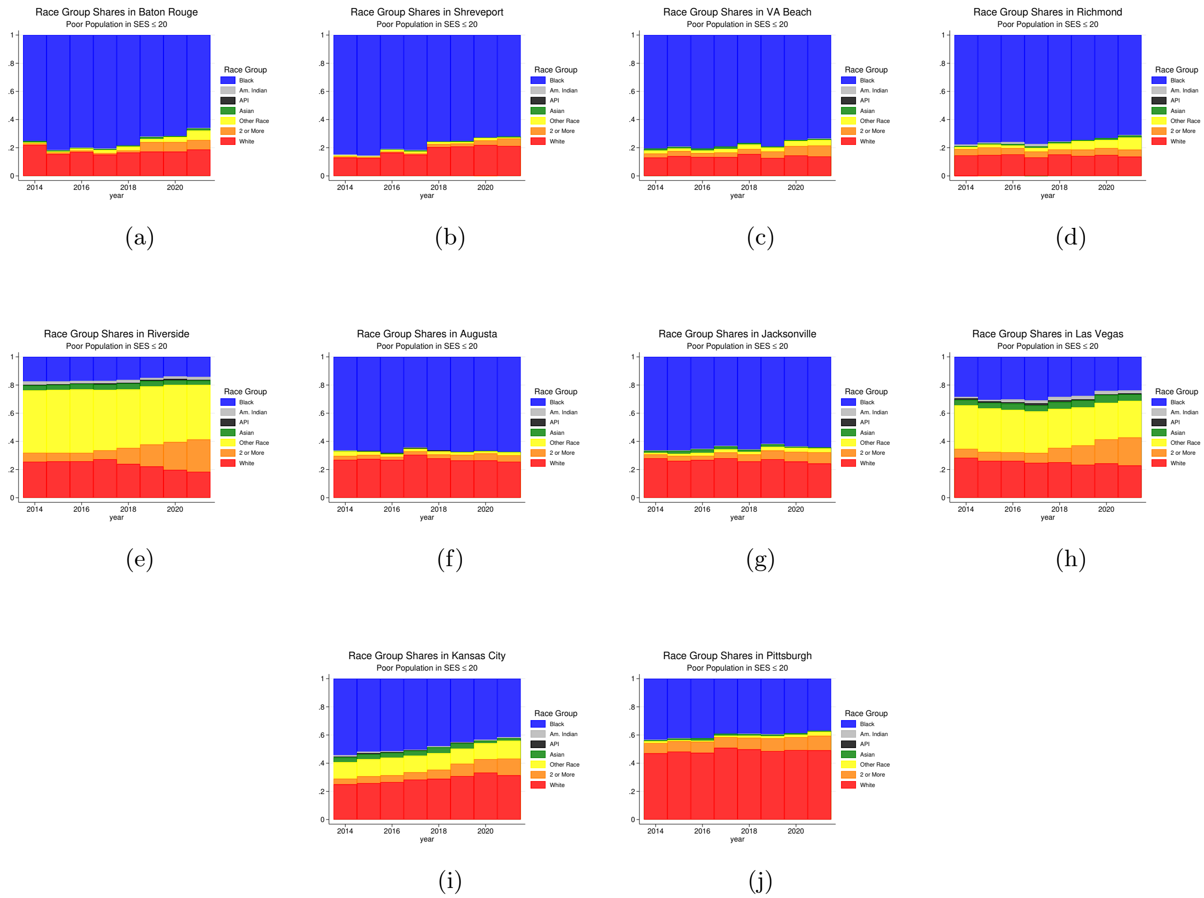


Figure 8: Racial Shares in the Bottom Quintile of Neighborhood SES, by City and Vintage of the American Community Survey (ACS)

# D Appendix: Black Population

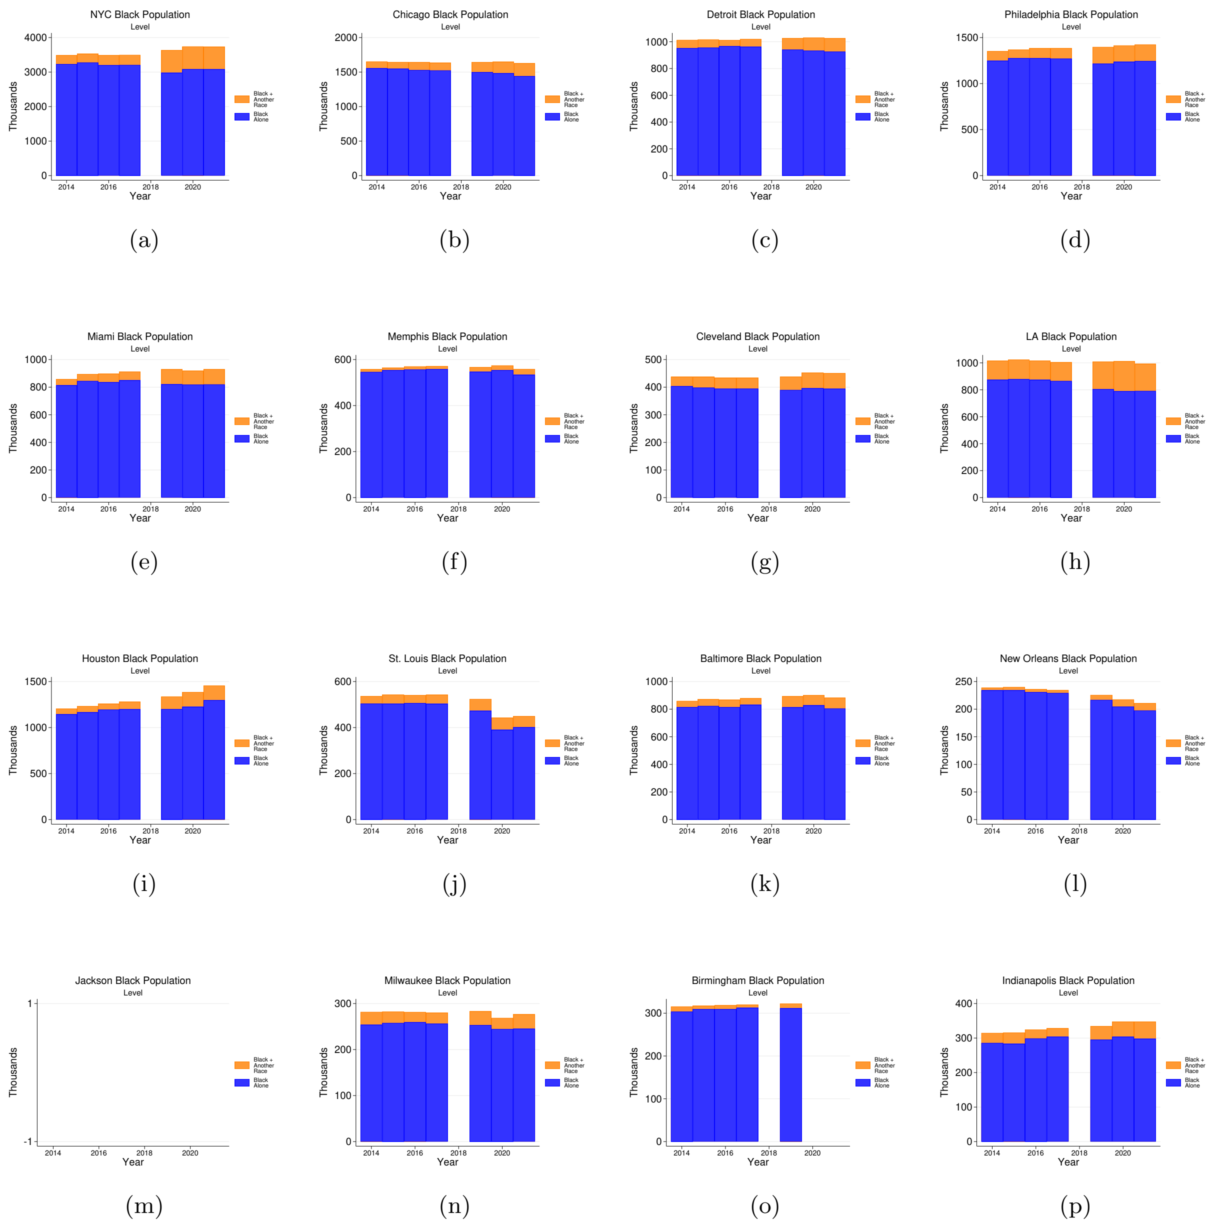


Figure 9: Black Population in the "Black Alone" or "Black and Another Race" Bins, by Vintage of the ACS

Note: The microdata examined in this figure are 1-year ACS estimates obtained from Ruggles et al. (2025). These calculations are based off of the detailed race variable with 100 bins. The "Black and Another Race" bin displayed here includes any of the 100 detailed race bins where "Black" is included.

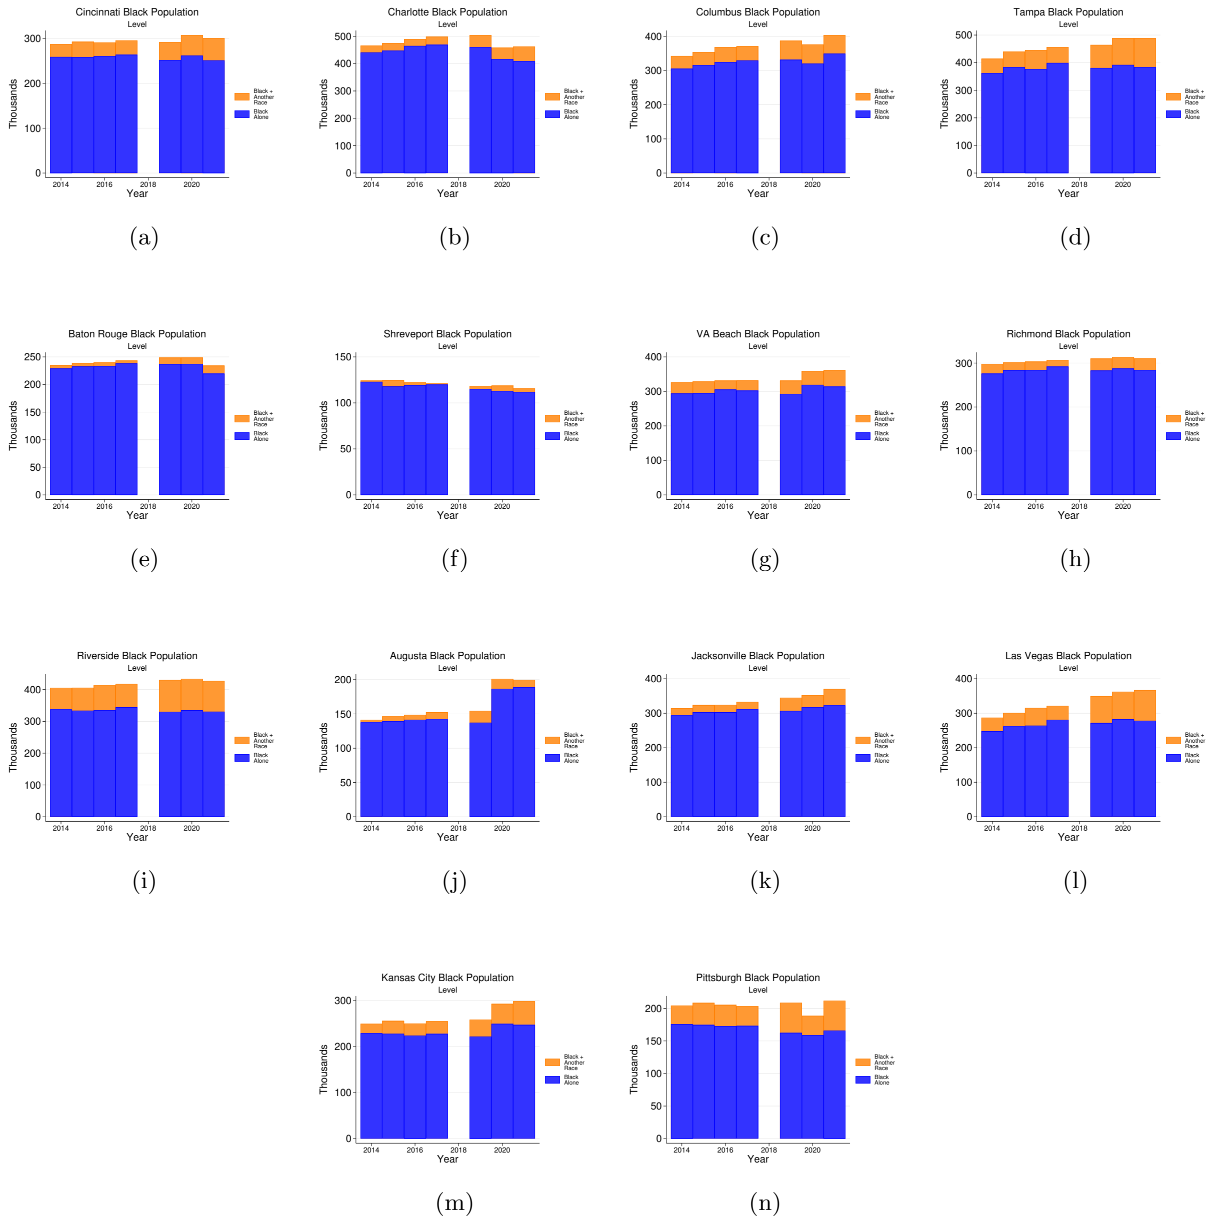


Figure 10: Black Population in the "Black Alone" or "Black and Another Race" Bins, by Vintage of the ACS

Note: The microdata examined in this figure are 1-year ACS estimates obtained from Ruggles et al. (2025). These calculations are based off of the detailed race variable with 100 bins. The "Black and Another Race" bin displayed here includes any of the 100 detailed race bins where "Black" is included.

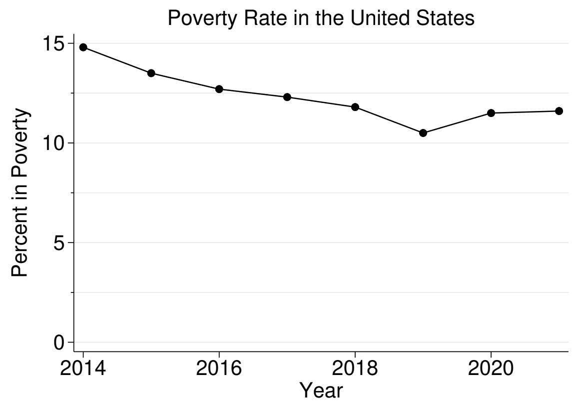


Figure 11: National Poverty Rate in the United States

Note: Data are obtained from Table A-3 of Shrider and Bijou (2025).

## E Appendix: Tract Boundary Changes

This appendix presents the full details behind the tract-boundary-change analysis in Section 3.1, culminating in Figure 3.

**Data.** At each decennial Census, tract boundaries are redrawn so that tract populations fall within the Census Bureau's target range. Our sample spans the 2015-2019 ACS (recorded on 2010-vintage tract boundaries) and the 2019-2023 ACS (recorded on 2020-vintage boundaries), so any comparison across these two vintages that does not adjust for the boundary revision, such as our Baseline estimates in Figure 2, risks confounding genuine demographic change with a pure artifact of redrawn tract lines.

**Robust Measure 1 (main-text technique): Area-Weighted reallocation.** Following Glaeser et al. (2025)'s construction of "reverse" versions of the Longitudinal Tract Database, we build a boundary-neutral counterpart to our headline series by area-weighting every 2019-2023 tract's data back onto the fixed set of 2010-vintage tracts, using the Census Bureau's official 2010-to-2020 Census Tract Relationship File, which records the area of geographic overlap between every pair of overlapping 2010- and 2020-vintage tracts.[7] For each 2020-vintage source tract $s$ overlapping a 2010-vintage target tract $t$, we compute the area share $w_{st} = \text{area}(s \cap t)/\text{area}(s)$ and reallocate $w_{st}$ of $s$'s poor Black population count onto $t$ – a standard area-weighted interpolation for a count variable. We similarly reallocate each source tract's neighborhood SES percentile onto every target tract it overlaps, but weighted instead by the *target* tract's own area share, $w'_{st} = \text{area}(s \cap t)/\text{area}(t)$, so that a target tract's reallocated SES quality is a population-agnostic area-weighted average of the quality of the source tracts covering it – the standard interpolation for a rate or quality variable, as opposed to a count. No tract's population is ever dropped in this procedure, and because 2015-2019 is already on 2010-vintage boundaries, the resulting series has no discontinuity at the 2020 boundary revision. We refer to the resulting boundary-neutral series as the area-weighted (AW) measure, and to the uncorrected, each-vintage-on-its-own-boundary series used elsewhere in the paper as the Baseline (R) measure.

**Target-Density-Weighted reallocations.** The area-weighted reallocation above implicitly assumes population is spread uniformly across each source tract's entire area – parks, water, and non-residential land included. We build two additional boundary-neutral series that are based on a different interpolation method that can weaken this assumption: target-density weighting (TDW; Schroeder (2007)), the method underlying the geographic crosswalks published by IPUMS NHGIS.[8] TDW weights each source-to-target tract overlap not just by its area, but by area times an ancillary population-density measure for the target tract. The result is that overlaps landing in denser target tracts receive more of the source tract's population than equally-sized overlaps landing in less-dense

[7] https://www2.census.gov/geo/docs/maps-data/data/rel2020/tract/tab20_tract20_tract10_natl.txt
[8] https://www.nhgis.org/geographic-crosswalks

ones:

$$w_{TDW}(s,t) = \frac{\text{area}(s \cap t) \times \text{density}_t}{\sum_{t'} \text{area}(s \cap t') \times \text{density}_{t'}}, \qquad \text{density}_t = \frac{\text{2020 (block-level) Black population of } t}{\text{land area of } t}.$$

We use two techniques for TDW:

**Robust Measure 2 (main-text technique): NHGIS's crosswalk weights.** In the main text, we implement TDW using the NHGIS's published crosswalk between 2020-vintage and 2012-vintage Census tracts. Rather than constructing a density measure and computing $w_{TDW}(s,t)$ ourselves from the formula above, we take the crosswalk's own `wt_pop` field directly as $w_{TDW}(s,t)$: NHGIS's expected share of source tract $s$'s total population located in target tract $t$, already computed by NHGIS from its 2020-2010 census block to census block crosswalk that was created using the same TDW methodology. For each 2020-vintage year, we multiply each source tract's poor Black population and neighborhood SES quality by `wt_pop` and sum by target tract, exactly as in the area-weighted reallocation above, then apply the bottom-decile threshold to the resulting target-tract quality estimate. This technique has two advantages over a custom-built alternative: it relies on a crosswalk that any researcher can obtain and apply without building their own block-level pipeline. And because the weight is population-based rather than race-specific, it does not mix the very 2020 race-processing methodology under study in this paper into the target-density measure itself. Its limitation is that it weights by total population rather than by Black population specifically, so it does not privilege the racial composition of the target tract the way a race-specific density measure would. We refer to the resulting boundary-neutral series as the target-density-weighted (TDW) measure, and use it throughout the main text and the rest of this appendix.

**Robust Measure 3: A custom, race-specific TDW technique.** As a further robustness check, we also implement TDW using each 2010-vintage target tract's actual 2020 Census block-level Black population as the density measure, obtained from the NHGIS (Schroeder et al. (2026); Table P2, 2020 Census P.L. 94-171 Redistricting Data Summary File), matching the NHGIS's published-crosswalk methodology but computing the weight ourselves with a race-specific rather than population-based density measure. Since 2020-vintage Census blocks nest within 2020-vintage tract boundaries rather than the 2010-vintage boundaries we need for the target geography, we assign each block to its containing 2010-vintage tract via a point-in-polygon join of the block's Census-reported internal point against 2010 TIGER/Line tract boundaries, and sum each target tract's block-level Black population directly.[9] For the small number of source tracts whose overlapping target tracts all have zero 2020 block-level Black population (157 of 27,618, or 0.6 percent), we fall back to the pure area weight.

[9] Because the density measure is thus drawn from the same 2020 race-processing methodology under study in this paper (unlike an earlier version of this check that used each target tract's own 2015-2019, pre-revision Black population as a tract-level proxy), we note that it enters TDW's weights only as a *relative* measure of density across a source tract's overlapping target tracts; any level shift from relabeling would move all of those target tracts' density estimates in roughly the same direction, so we do not think this compromises the check.

This custom technique's own headcount bias statistics are close to the main-text technique's: a median bias of −3.4 percent of a city's 2017 population (mean −3.9 percent), against −2.3 percent (mean −2.8 percent) for the main-text technique, with Houston again the largest outlier (−14.4 percent, against −12.8 percent for the main-text technique). Figure 12 shows the resulting Baseline-vs-custom-TDW slope chart, the direct analogue of Figure 14a below. We refer to the resulting change in concentrated poverty under the main-text (NHGIS crosswalk) technique as **Measure 2** and under the custom block-level technique as **Measure 3**. Figure 13 compares Measure 2 and Measure 3 directly: the two are correlated at 0.99 across the 30 cities, with a median absolute difference of 0.3 percentage points (largest: Augusta, at 4.5 percentage points). Because the two techniques agree this closely, we use the main-text (NHGIS crosswalk) technique throughout the remainder of this appendix.

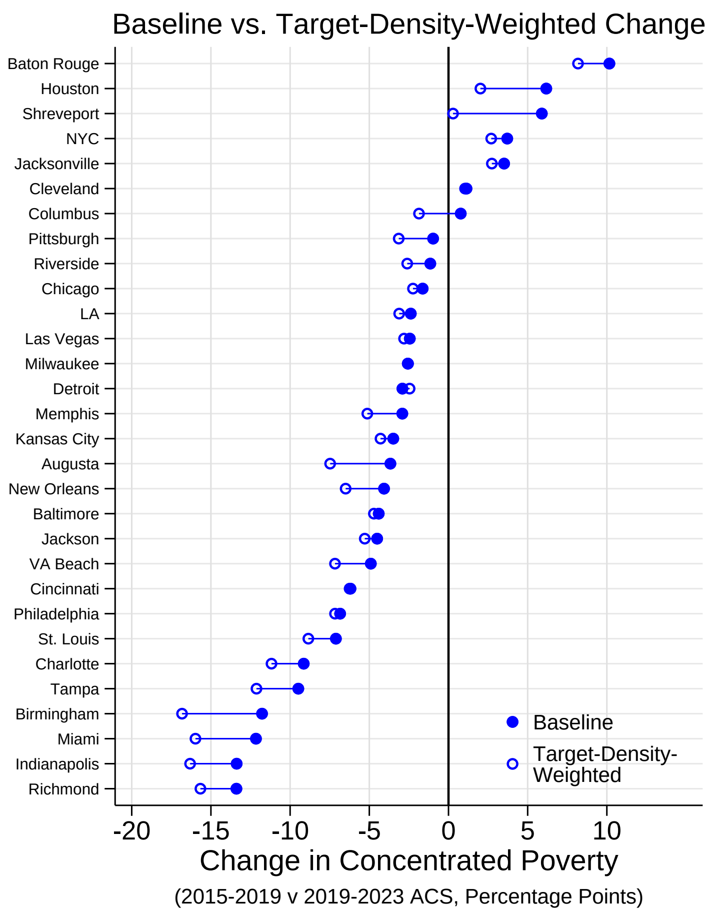


Figure 12: Baseline vs. Custom Race-Specific TDW Change in Concentrated Poverty, by City

Note: Each row is one of the 30 cities in our sample, sorted from the most negative to the most positive Baseline measure of Concentrated Poverty, using the same city ordering as Figure 14a below. The solid blue circle represents the Baseline measure; the hollow blue circle represents the custom block-level Target-Density-Weighted measure.

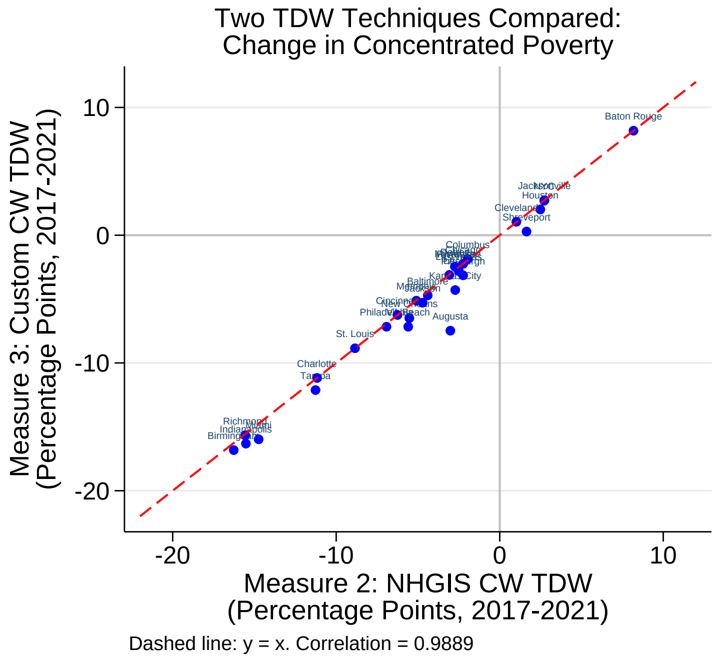


Figure 13: Two TDW Techniques Compared: Change in Concentrated Poverty, across 30 Cities

Note: Each point is one of the 30 cities in our sample. Horizontal axis: Measure 2, the change in concentrated poverty under the main-text (NHGIS crosswalk) TDW technique. Vertical axis: Measure 3, the analogous change under the custom block-level TDW technique. The dashed line is $y = x$. Correlation across the 30 cities is 0.9889.

**A direct test: is the boundary revision driving the measured decline?** Let $B(y)$ denote the Baseline total number of poor Black residents in bottom-decile-SES ($\leq$ 10th percentile) tracts in year $y$, using each tract's own vintage boundary and own SES ranking. This is the measure we use throughout the analysis, with no boundary correction. Let $AW(y)$ denote the corresponding Area-Weighted total from the reallocation above, and $TDW(y)$ the analogous total under Target-Density Weighting, described above. Because $B$, $AW$, and $TDW$ only begin to diverge with the 2018-2022 ACS, the entire contribution of the boundary revision to the 2021 total is the single quantity $AW(2021) - B(2021)$ (or $TDW(2021) - B(2021)$). We scale this by each city's 2017 population, $B(2017)$ since this is the last pre-revision year. This population is therefore the natural reference for a bias that only arises with the revision of tract boundaries rather than by the measured decline itself.

Table 1 reports $B(2017)$, $B(2021)$, both reallocated totals ($AW(2021)$ and $TDW(2021)$), and both methods' bias statistics side by side, for all 30 cities in our sample. Three features stand out. First, both biases are uniformly small: the median is $-2.4$ percent of a city's 2017 population under area-weighting and $-2.3$ percent under TDW; the means are $-3.3$ and $-2.8$ percent, respectively. Second, both are essentially always negative: area-weighting is negative in 28 of the 30 cities (all but Detroit and Milwaukee, each a modest positive value under half a percentage point), and TDW is negative in 25 of 30 (Detroit, Baltimore, Milwaukee, Augusta, and Kansas City the exceptions, each a modest positive value, the largest being Kansas City at 1.8 percent), meaning the boundary-neutral decline is, if anything, *larger* than the Baseline, as-measured decline under either method. Third, Houston is the largest outlier under both methods, at $-12.8$ percent under area-weighting and $-12.8$ percent under TDW, it represents the single largest bias in the table. Where the tract-boundary revision has any systematic effect at all, under either interpolation method, it *attenuates* the measured decline rather than manufacturing or inflating the appearance of one. If the revision were instead inflating the appearance of a decline, we would expect a positive bias in at least some cities; area-weighting shows only two such cities, both close to zero, while TDW shows five, three of which (Baltimore, Milwaukee, Detroit) are close to zero and two (Augusta, Kansas City) more modest.

Table 1: Boundary-Redefinition Bias in the 2021 Total, by City

| | Count | | | | Bias | | | |
|---|---|---|---|---|---|---|---|---|
| City | $B$ <br> $B(2017)$ | $B'$ <br> $B(2021)$ | $AW'$ <br> $AW(2021)$ | $TDW'$ <br> $TDW(2021)$ | $B' - AW'$ | $B' - TDW'$ | $[B' - AW']$ <br> $/B$ | $[B' - TDW']$ <br> $/B$ |
| NYC | 266.8 | 272.5 | 266.1 | 266.8 | -6.34 | -5.70 | -0.02 | -0.02 |
| Chicago | 189.9 | 169.8 | 167.6 | 167.7 | -2.25 | -2.09 | -0.01 | -0.01 |
| Detroit | 180.8 | 153.3 | 154.1 | 153.7 | 0.76 | 0.40 | 0.00 | 0.00 |
| Philadelphia | 157.3 | 124.1 | 123.9 | 123.9 | -0.16 | -0.21 | 0.00 | 0.00 |
| Miami | 128.5 | 80.5 | 75.0 | 74.8 | -5.58 | -5.69 | -0.04 | -0.04 |
| Memphis | 91.8 | 79.8 | 76.7 | 76.7 | -3.10 | -3.15 | -0.03 | -0.03 |
| Cleveland | 80.1 | 72.8 | 72.5 | 72.6 | -0.25 | -0.13 | 0.00 | 0.00 |
| LA | 70.7 | 59.4 | 59.3 | 58.3 | -0.10 | -1.09 | 0.00 | -0.02 |
| Houston | 64.5 | 83.0 | 74.7 | 74.7 | -8.26 | -8.26 | -0.13 | -0.13 |
| St. Louis | 60.9 | 45.1 | 43.3 | 43.2 | -1.82 | -1.86 | -0.03 | -0.03 |
| Baltimore | 58.9 | 49.2 | 48.8 | 49.2 | -0.43 | 0.00 | -0.01 | 0.00 |
| New Orleans | 57.4 | 44.7 | 42.8 | 43.3 | -1.97 | -1.42 | -0.03 | -0.02 |
| Jackson | 47.2 | 41.5 | 40.6 | 41.3 | -0.83 | -0.15 | -0.02 | 0.00 |
| Milwaukee | 45.7 | 38.7 | 38.7 | 38.7 | 0.04 | 0.00 | 0.00 | 0.00 |
| Birmingham | 41.8 | 33.3 | 29.6 | 30.0 | -3.69 | -3.30 | -0.09 | -0.08 |
| Indianapolis | 34.3 | 22.8 | 19.9 | 21.4 | -2.87 | -1.40 | -0.08 | -0.04 |
| Cincinnati | 34.0 | 27.4 | 27.3 | 27.4 | -0.06 | -0.04 | 0.00 | 0.00 |
| Charlotte | 33.6 | 23.2 | 21.3 | 21.4 | -1.90 | -1.88 | -0.06 | -0.06 |
| Columbus | 32.7 | 30.4 | 28.2 | 28.3 | -2.16 | -2.08 | -0.07 | -0.06 |
| Tampa | 32.4 | 21.1 | 19.8 | 19.8 | -1.27 | -1.25 | -0.04 | -0.04 |
| Baton Rouge | 32.1 | 36.9 | 34.7 | 35.5 | -2.22 | -1.34 | -0.07 | -0.04 |
| Shreveport | 30.9 | 30.6 | 28.4 | 28.6 | -2.17 | -2.00 | -0.07 | -0.06 |
| VA Beach | 29.9 | 24.6 | 24.0 | 24.0 | -0.60 | -0.65 | -0.02 | -0.02 |
| Richmond | 26.3 | 15.6 | 14.2 | 14.3 | -1.45 | -1.31 | -0.05 | -0.05 |
| Riverside | 26.1 | 23.9 | 23.0 | 23.0 | -0.92 | -0.86 | -0.04 | -0.03 |
| Augusta | 26.0 | 22.6 | 22.3 | 22.9 | -0.33 | 0.32 | -0.01 | 0.01 |
| Jacksonville | 26.0 | 29.0 | 28.4 | 28.4 | -0.56 | -0.55 | -0.02 | -0.02 |
| Las Vegas | 24.9 | 22.7 | 22.6 | 22.6 | -0.06 | -0.03 | 0.00 | 0.00 |
| Kansas City | 22.6 | 18.7 | 18.2 | 19.1 | -0.55 | 0.41 | -0.02 | 0.02 |

*continued on next page*

*Table 1 continued*

| City | Count | | | | Bias | | | |
|---|---|---|---|---|---|---|---|---|
| | $B$ | $B'$ | $AW'$ | $TDW'$ | | | $[B'-AW']$ | $[B'-TDW']$ |
| | $B(2017)$ | $B(2021)$ | $AW(2021)$ | $TDW(2021)$ | $B'-AW'$ | $B'-TDW'$ | $/B$ | $/B$ |
| Pittsburgh | 18.3 | 18.0 | 17.9 | 17.3 | -0.02 | -0.67 | 0.00 | -0.04 |
| Median | | | | | | | -0.02 | -0.02 |
| Mean | | | | | | | -0.03 | -0.03 |

Table Note: City is ordered by 2015-2019 ACS poor Black population in bottom-decile-SES tracts. $B(2017)$ and $B(2021)$ are the Baseline totals of poor Black residents in bottom-decile-SES tracts, using each year's own tract vintage and own SES ranking, in thousands. $B(2017)$ is the last pre-revision year (ie, the 2015-2019 ACS). $AW(2021)$ is the area-weighted reallocation of the 2021 total onto fixed 2010-vintage tract boundaries; $TDW(2021)$ is the analogous target-density-weighted reallocation (both methods are defined in full below). $B'-AW'$ and $B'-TDW'$ are $AW(2021)-B(2021)$ and $TDW(2021)-B(2021)$, in thousands (here $B' \equiv B(2021)$ and $AW' \equiv AW(2021)$, $TDW' \equiv TDW(2021)$). The last two columns express these as a share of the city's 2017 population, $B(2017)$. A negative bias means the boundary-neutral decline is at least as large as the Baseline, so that the boundary revision is not inflating the appearance of decline.

Figures 14a and 14c isolate the Baseline-vs-Measure-2 (NHGIS crosswalk TDW) and Baseline-vs-Area-Weighted comparisons, respectively, from Figure 3 in the main text, so each reallocated measure can be compared against the Baseline on its own; Figure 14b adds the analogous comparison for Measure 3 (custom block-level TDW), which does not appear in the main-text figure. All three patterns are visually almost identical – the clearest illustration of how little the reallocation's conclusion depends on the choice of interpolation method – and, in all three cases, accounting for the tract-boundary revision makes the decline in concentrated poverty larger, not smaller, than the Baseline measure.

Finally, we note that the negative bias documented above is not a mechanical by-product of the reweighting procedures implemented. We considered several simulated tract splits and find that the reallocation methods can produce both positive and negative changes relative to the baseline.

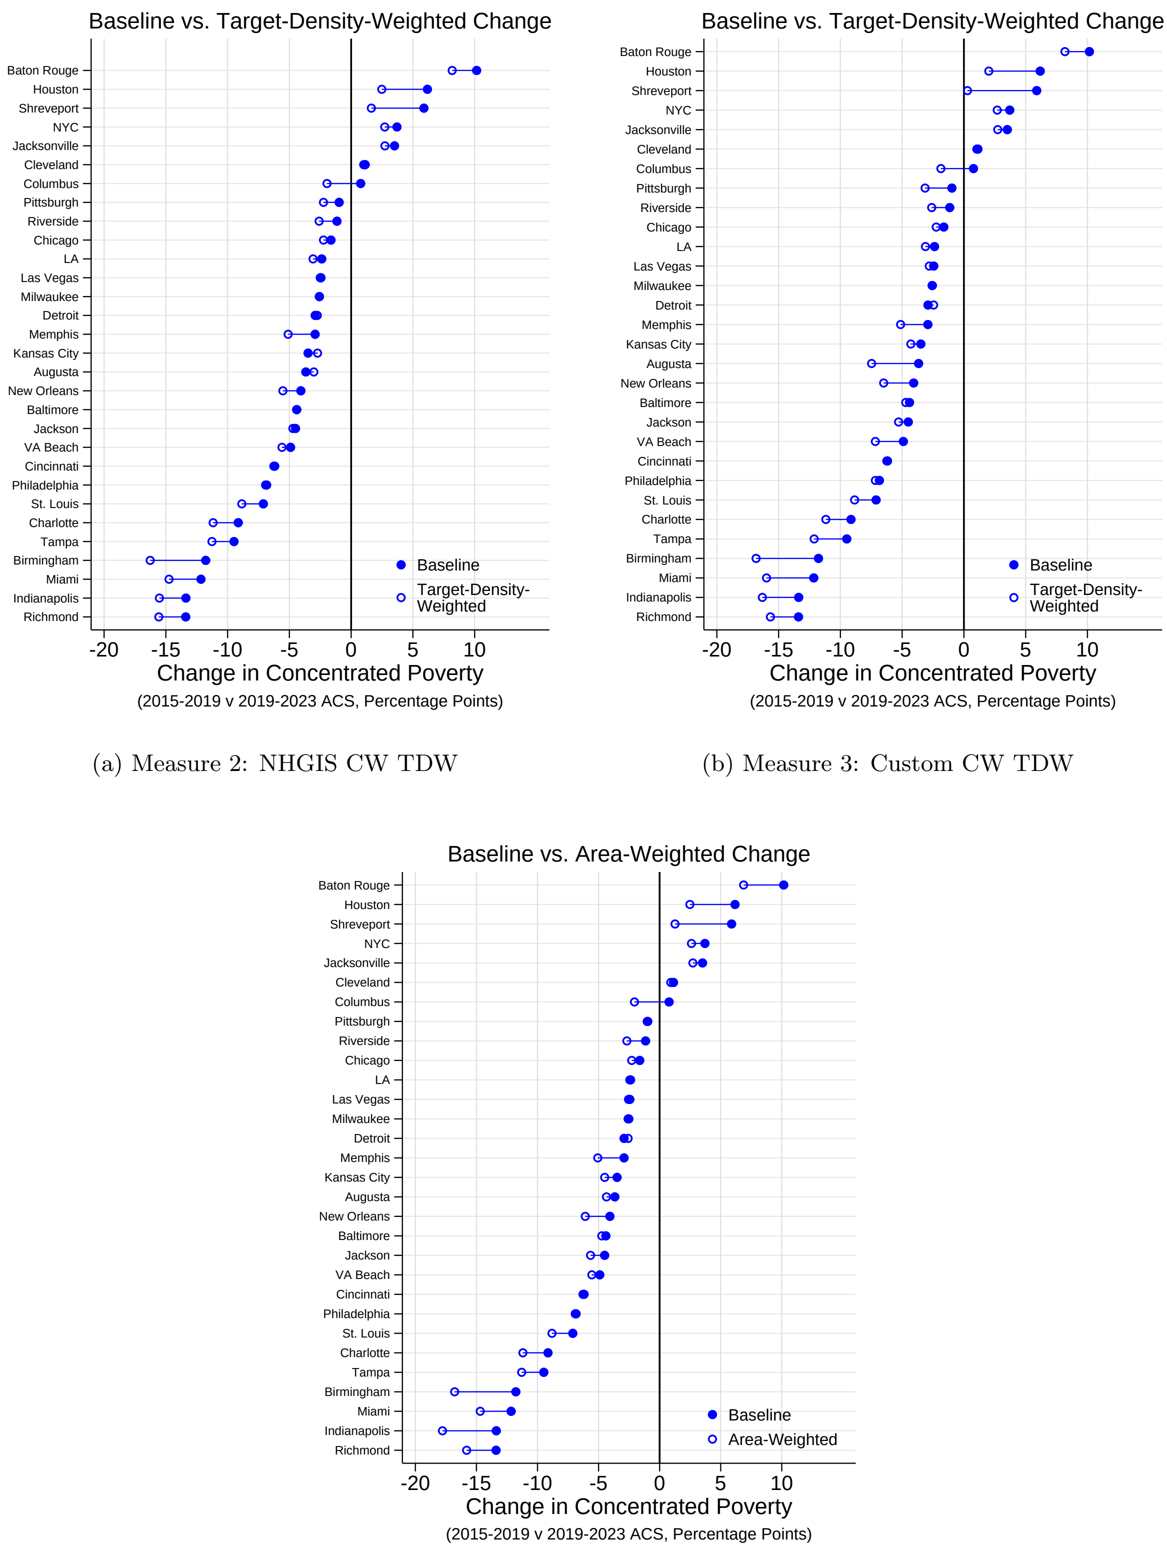


(a) Measure 2: NHGIS CW TDW

(b) Measure 3: Custom CW TDW

(c) Area-Weighted

Figure 14: Baseline vs. Reallocated Change in Concentrated Poverty, by City

Note: Each row is one of the 30 cities in our sample, sorted from the most negative to the most positive Baseline measure of Concentrated Poverty. The solid blue circle represents the Baseline measure; the hollow blue circle represents Measure 2 (NHGIS crosswalk TDW, panel a), Measure 3 (custom block-level TDW, panel b), or Measure 1 (Area-Weighted, panel c).