sociological science

The Causal Impact of Segregation on a Disparity: A Gap-Closing Approach

Ian Lundberg

UCLA

Abstract: Segregation—whether across schools, neighborhoods, or occupations—is regularly invoked as a cause of social and economic disparities. However, segregation is a complicated causal treatment: what do we mean when we appeal to a world in which segregation does not exist? One could take societal contexts as the unit of analysis and compare across societies with differing levels of segregation. In practice, it is more common for studies of segregation to take persons or households as the unit of analysis within a single societal context, focusing on what would happen if particular individuals were counterfactually assigned to social positions in a more equitable way. Taking this latter framework, this article shows how to study segregation as a cause. The first step is to theorize a counterfactual assignment rule: what would it mean to assign people to social positions equitably? The second step is to identify the causal effect of those social positions and simulate counterfactual outcomes. The third step is to interpret results as the impact of a unit-level (rather than society-level) intervention. A running example and empirical analysis illustrates the approach by studying the causal effect of occupational segregation on a racial health gap.

Keywords: segregation; causal inference; inequality; quantitative methods; health disparities

Reproducibility Package: doi.org/10.7910/DVN/TB5Q4N.

Citation: Lundberg, Ian. 2025. "The Causal Impact of Segregation on a Disparity: A Gap-Closing Approach" Sociological Science 12: 871-890.

Received: July 15, 2025

Accepted: August 31, 2025

Published: December 9, 2025

Editor(s): Arnout van de Rijt, Maria Abascal

DOI: 10.15195/v12.a35

SEGREGATION is a descriptive fact in unequal societies, and it also plays a causal role in generating further inequality. As one example, racial segregation of households in the United States is a causal force that contributes to racial disparities in wealth through unequal appreciation of home values (Flippen 2004; Markley et al. 2020), disparities in children's access to well-funded schools (Weathers and Sosina 2022), and disparities in families' exposure to environmental toxins (Sampson and Winter, 2016). As an example with a different kind of segregation, sex segregation of people across occupations is often framed as a key contributor to the sex gap in pay (Blau and Kahn 2017; Petersen and Morgan 1995). From a theoretical perspective, it is conceptually plausible that disparities would narrow in a counterfactual world in which various forms of segregation were eliminated.

But how should one study that sort of counterfactual empirically? In practice, many researchers analyze data on only one society. The researcher is then only able to observe that society under one pattern of segregation: whatever pattern factually exists. However, the researcher still has access to variation across persons (or households) who are unequally allocated to social positions, which is the individual-level process that aggregates to societal-level segregation. This article shows how to use that empirical variation to study the causal impact of segregation. The first step is to define what it would mean to allocate people to social positions equitably. The second step is to causally study the counterfactual outcomes that would be realized

under that assignment rule. The third step is to interpret the results locally, aware that the analysis was carried out with an individual (instead of a society) as the unit of analysis.

The approach of this article builds on methodological research that decomposes disparities according to their causal inputs (Jackson and VanderWeele 2018; Lundberg 2024). The methodological contribution of the present article is how to carry out these decompositions when the causal input of interest is categorical—occupations, neighborhoods, and schools—and when this input may be substantially shaped by measured confounders. Instead of theorizing a world where units were assigned to a counterfactual treatment value, this article theorizes a world where units are assigned to social positions by a counterfactual assignment rule. The broader contribution of the present article is to show in a sociological setting how to use variation and causal identification at the level of individuals to study the causal impact of an input (segregation) that is more often defined at the level of a society.

Motivating Example: Occupational Segregation May Contribute to Racial Health Disparities

In 2019, 2.8 million Americans experienced workplace injuries (Bureau of Labor Statistics 2020). Compared with white workers, black workers disproportionately hold occupations with high injury rates (Seabury, Terp, and Boden 2017; Stanbury and Rosenman 2014). Case studies reveal specific links between particular occupational hazards and health outcomes, from heat-related illness among farmworkers (Fleischer et al. 2013) to physical trauma among day laborers (Lowry et al. 2010) and hotel workers (Buchanan et al. 2010). Occupations also expose individuals to unequal toxins with effects that may accumulative over time. For example, Juon et al. (2021) find that black workers hold occupations that place them at higher exposure to dangers such as silica and asbestos, which increase the risk of lung cancer. Beyond direct injuries, a person's occupation can indirectly affect health through pathways such as work hours and scheduling that shape sleep and exercise (Harknett, Schneider, and Wolfe 2020; Moen et al. 2011), or through changes in other risk behaviors such as alcohol consumption (Prins et al. 2019). There are numerous reasons to expect both direct and indirect links between the workplace environment and health outcomes (Ahonen et al. 2018; Burgard and Lin 2013; Lipscomb et al. 2006; Schulte et al. 2017).

Occupational risks may contribute to racial health disparities because, despite the Civil Rights Act of 1964, occupations remain extensively segregated by race (Ferguson and Koning 2018; Stainback and Tomaskovic-Devey 2012; Tomaskovic-Devey et al. 2006). Figure 1 presents a descriptive summary using data from the Current Population Survey (described in detail in a later section). Each occupation (represented by a dot) is scored on the x-axis by racial composition. On the y-axis, each occupation is scored by the rate of onset of work-limiting disability: the proportion of non-disabled workers in year t to report in year $t + 1$ that they have a disability limiting the type or amount of work they can do. The descriptive relationship between these variables is strong. Occupations that are disproportionately held by non-Hispanic black or Hispanic workers tend to have high rates of onset of

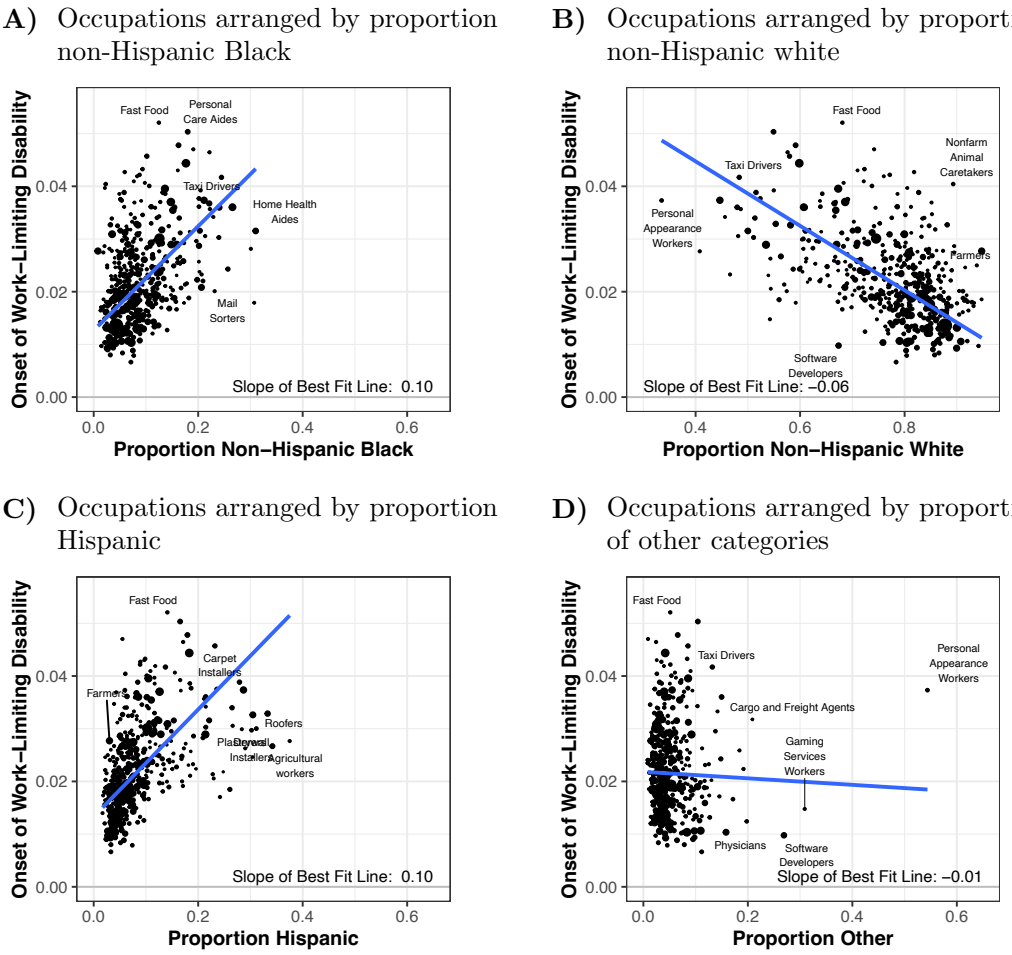

Figure 1: Dangerous occupations are disproportionately filled by people of color. Each point is an occupation scaled proportional to its size in the full population. The horizontal dimension is an estimate of the proportion identifying with a given racial and ethnic category in the occupation in year t. The vertical dimension is the proportion of all people in that occupation for whom a disability or health condition limiting the kind or amount of work they can do is reported in year $t + 1$. Both the horizontal and vertical dimensions are estimated by multilevel models with random intercepts using data from the 2005 to 2020 Annual Social and Economic Supplement of the Current Population Survey (details in the online supplement B).

work-limiting disability: home health aides, agricultural workers, and roofers. Occupations that are disproportionately held by white or other workers tend to have low onset of work-limiting disability: software developers, physicians, and CEOs.

It is conceptually plausible that macro-level occupational segregation is a contributor to macro-level racial health disparities. One could carry out the analysis at the macro-level, taking the unit of analysis to be social contexts such as states or countries that vary in their macro-level characteristics (Heymann 2003; Ólafsdóttir and Beckfield 2020). For example, in a study of residential segregation, Torrats-Espinosa (2021) estimates the causal effect of county-level residential segregation on county-level racial disparities in COVID-19 mortality, net of control variables

defined at the county level. All of the anaysis is at the macro (county) level rather than the micro (person) level. Although macro-level analysis is powerful when data enable it, many researchers operate with data only at the micro-level within one macro-contex. For example, they seek to make a claim about one social context (e.g., the United States) using data on units within that social context (e.g., persons). To study macro-level causes using micro-level data requires careful reasoning about how individual-level evidence can aggregate up to produce macro-level patterns (Bearman, Moody, and Stovel 2004; Coleman 1990; Hedström and Bearman 2009). The next section illustrates how to study the impact of occupational segregation on a racial disparity using micro-level data in the potential outcomes framework of causal inference (Imbens and Rubin 2015).

Data for Empirical Illustration

To make the discussion below concrete, the analysis uses data from the 2005 to 2020 Current Population Survey (hereafter CPS) on the civilian U.S. population aged 25−60, harmonized by the Integrated Public Use Microdata Series (IPUMS; Flood et al. 2020). The outcome variable is whether each individual reports "a health problem or a disability which prevents him/her from working or which limits the kind or amount of work." I refer to this as a report of a work-limiting disability.[1] Occupation (the treatment) is measured in three-digit codes harmonized by IPUMS (Flood et al. 2020) to a standard based on the 2010 Census coding system, which contains 451 unique non-military occupations (e.g., chefs and cooks, elementary and middle school teachers, cashiers, and roofers). Racial category is self-reported (non-Hispanic black, non-Hispanic white, Hispanic, and other). The data also contain covariates such as sex (men and women), education (less than high school, high school, some college, and college), foreign born, age, year, and a lagged measure of the standard five-point scale of self-rated health.[2]

All predictor variables (race, occupation, and covariates) are measured at year t. The outcome is measured 12 months later, taking advantage of the panel structure of the CPS (online supplement Fig. 13, see Rivera Drew et al. 2014). The main analyses restrict the sample to those who report at year t that they have never quit a job for health reasons, are currently employed, and do not have a work-limiting disability (Fig. 2). This restriction on the lagged outcome and employment reduces the risk that differences in the outcome observed one year later are consequences of latent health. The analytical sample contains 244,224 persons.

How to Study the Impact of Segregation Using Person-Level Potential Outcomes

The empirical task involves two steps: make an estimate for each person and then aggregate over people. To define the causal estimand, I adopt potential outcomes notation (Imbens and Rubin 2015). Let $y_i(d)$ represents the health (healthy or with a work-limiting disability) that person i would realize if they held occupation d (possibly counter to fact). Because there are 451 occupations in the Census coding system, each person has 451 potential outcomes.[3] Yet we observe only one: the outcome in the occupation they factually held.

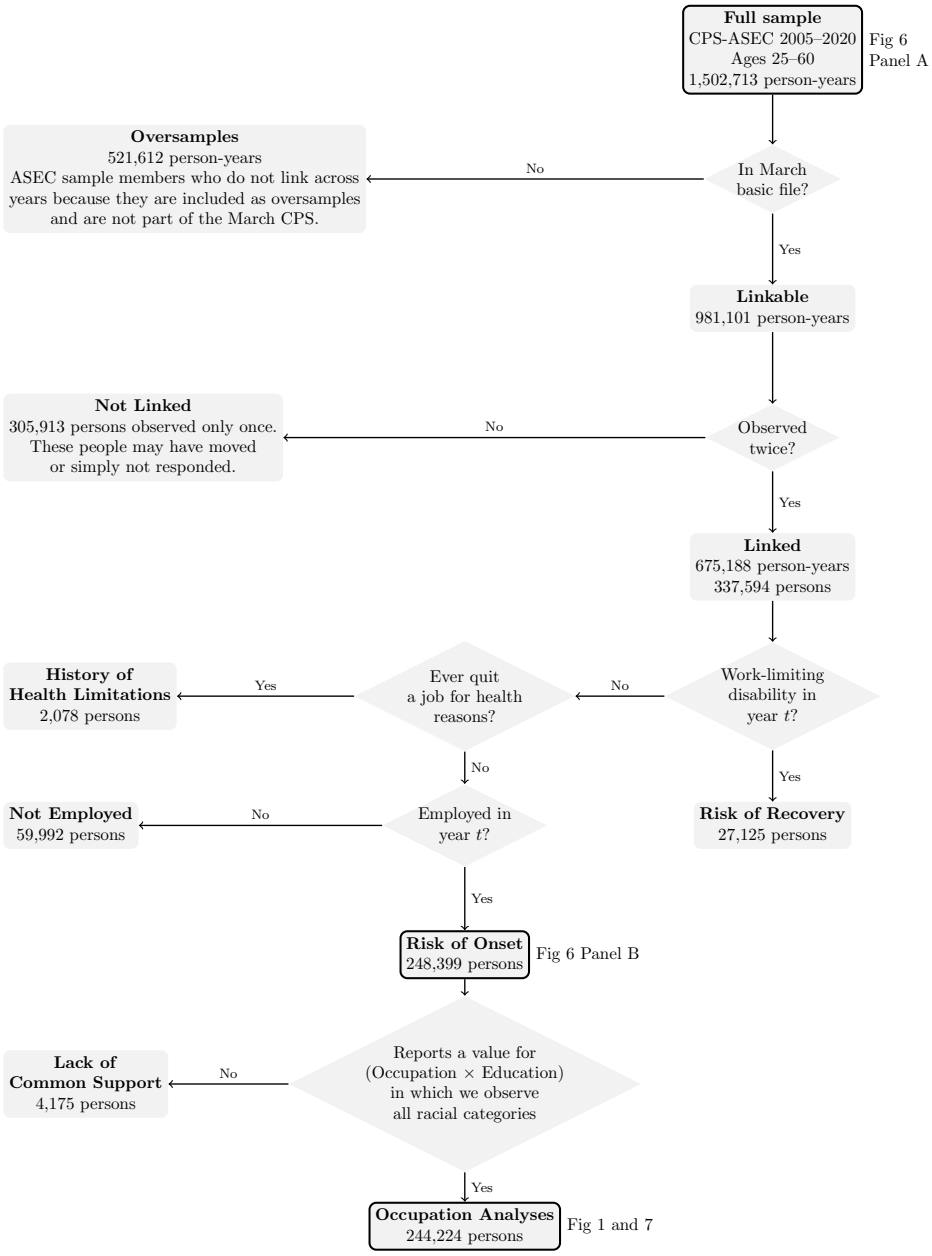

Figure 2: Sample restrictions. The three rectangles with black outlines correspond to samples analyzed in the main text.

To illustrate, consider a hypothetical population of eight people assigned to one of three occupations (Fig. 3). We are interested in the health disparity across two categories of people, here labeled abstractly as circles and diamonds but which could correspond to categories of race, class, or gender. In this hypothetical setting, Panel A shows a health disparity: only $\frac{1}{4}$ of the circles will develop a disability next year compared with $\frac{3}{4}$ of the diamonds. That disparity could be a consequence of segregation: the diamonds are disproportionately home health aids, an occupation

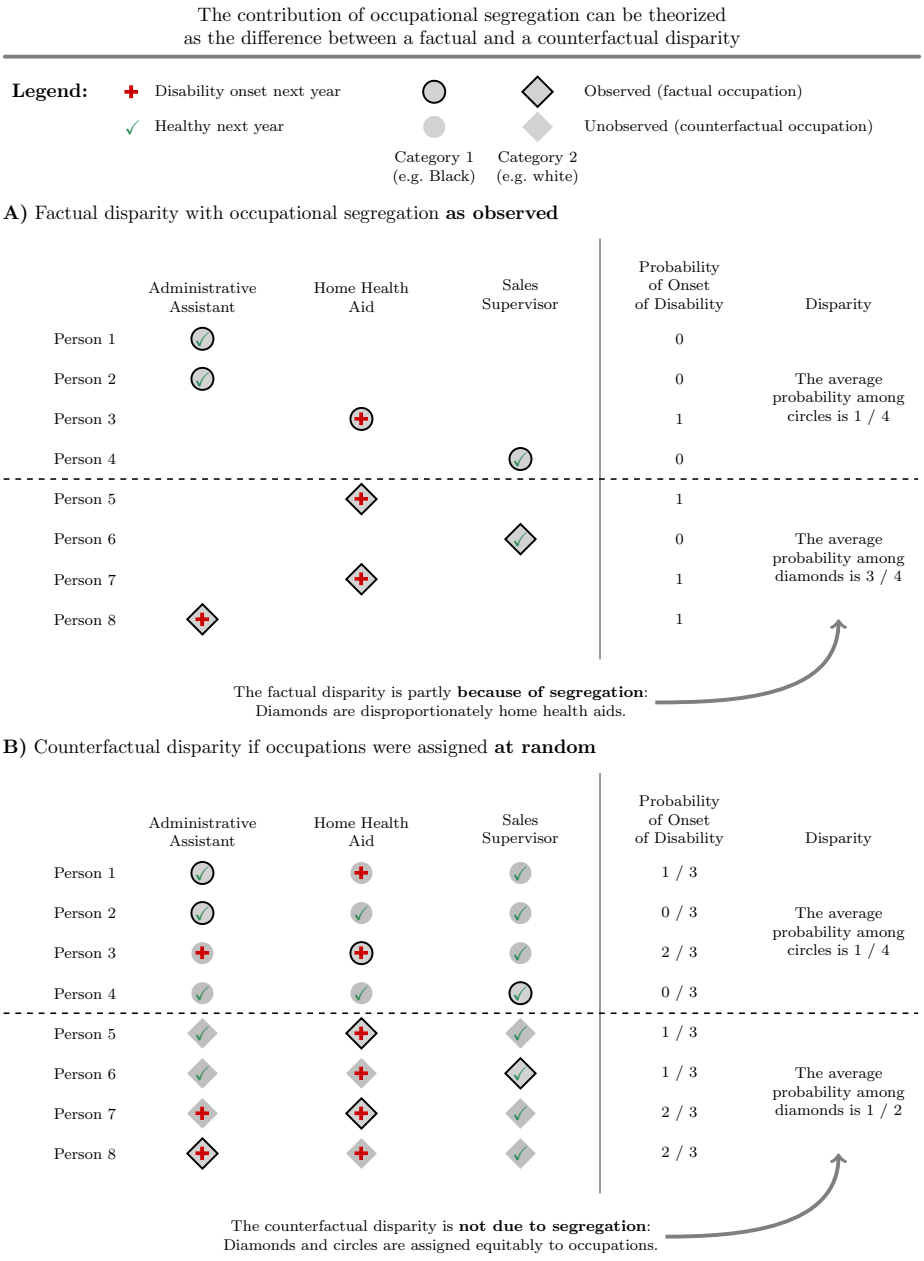

Figure 3: Consequences of segregation in the potential outcomes framework. Consider eight people assigned to three possible occupations. Panel A depicts occupations that occur, in which the factual health disparity is partly a consequence of occupational segregation. Panel B depicts a hypothetical setting where occupations are instead assigned at random with equal probabilities. If we were able to learn about the counterfactual setting, the difference between the factual and counterfactual settings would quantify the contribution of occupational segregation to the health disparity.

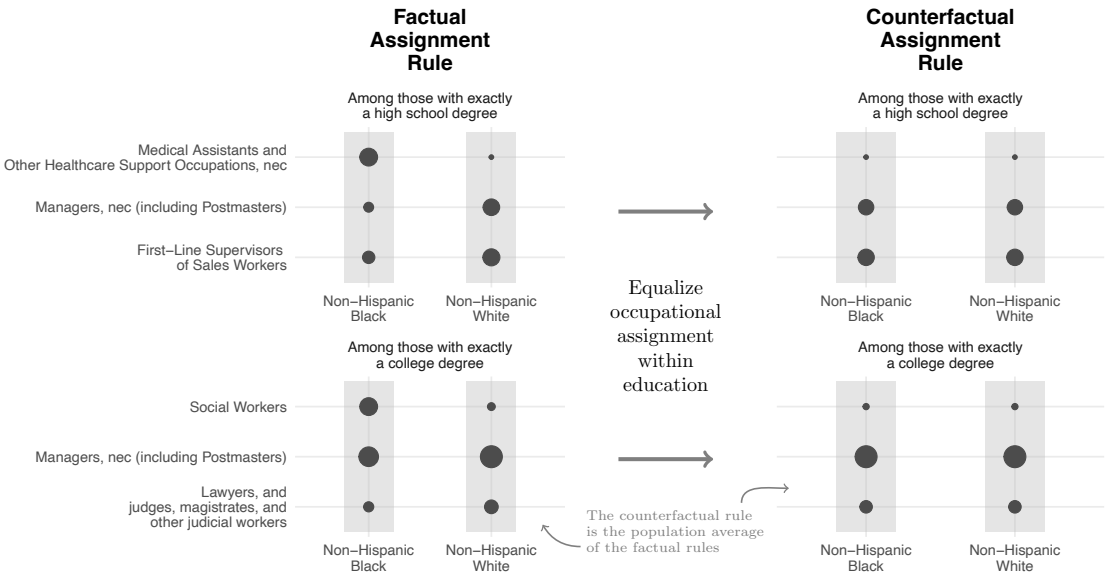

Figure 4: A counterfactual assignment rule illustrated with selected occupations. The size of each dot corresponds to the probability of assignment to that occupation. The factual assignment rule is how occupations are assigned in observational data, as a function of not only education but also race and other variables. In the counterfactual assignment rule, the probability of assignment to each occupation is the same across all subgroups (e.g., race) within populations defined by education.

which (in this hypothetical example) is very hazardous. What if each person was instead assigned to a random occupation, so that the disparity in that counterfactual world could not be attributed to occupational segregation? If we knew the full table of potential outcomes (Panel B), we could answer that question: randomly assigning occupations would reduce the disparity from $\frac{1}{2}$ to $\frac{1}{4}$. That provides a metric of the degree to which occupational segregation contributes to a disparity. Importantly, there is no notion of changing a circle into a diamond or vice versa; the causal effect of the category (circles and diamonds) remains undefined. The summary is about how much the gap between the circles and diamonds (two subpopulations of people) would close if counterfactually assigned to occupations differently (similar to the gap closing estimand from Lundberg 2024).

Theorize a Counterfactual Rule to Assign Individuals to Social Positions

The illustration above considered a counterfactual world where occupations were assigned with equal probabilities, such that each person has a 1/3 probability of assignment to each of the three possible occupations. However, this assignment rule is unrealistically simple for two reasons. First, some occupations are larger than others and should be assigned at higher probabilities. Second, some occupations are not accessible to some people, as discussed in other research that analyzes occupation as a causal treatment (Lundberg, Molitor, and Brand 2025). For example, someone who did not finish high school could not become a high school teacher. This constraint limits the scope of counterfactuals that can be empirically studied,

as we revisit later in a section on scope conditions. To study an intervention that is empirically tractable, one must define a more realistic assignment rule in which every unit's counterfactual treatment is one that the factual world might assign to that unit with non-zero probability.

This article defines a counterfactual assignment rule that desegregates occupations within educational categories,

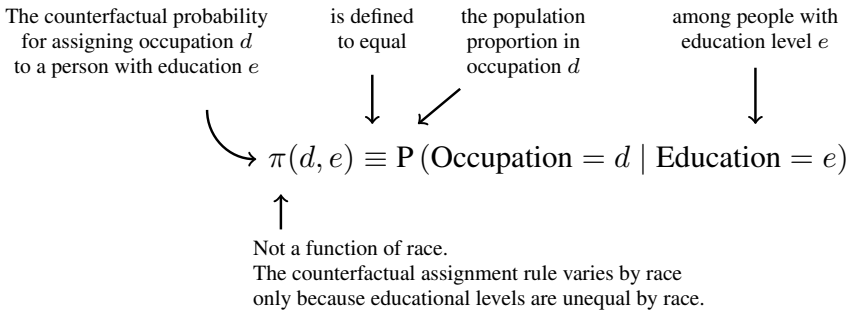

(1)

such that the probability of being assigned to occupation d for a person with education level e equals the proportion holding that occupation among all people with that education category. This assignment rule is counterfactual because it is different from how occupations are factually assigned: in the real world, the probability of holding a particular occupation is a function of not only education but also race, sex, age, et cetera (illustrated in Fig. 4).

The choice of a counterfactual assignment rule is non-trivial. It is important to recognize that a world with occupations assigned by Eq. 1 would still have some racial segregation: to the degree that black respondents have less education than white respondents, for example, they would be assigned to different occupations. When defining a counterfactual assignment rule, one should think carefully about both an assignment rule that is realistic (e.g., someone who did not finish high school could not teach high school) and also equitable and socially just (e.g., a world where occupations are assigned as a function of sex might be considered unjust, see Jackson 2021).

Identify by Causal Assumptions

To estimate the racial gap under a counterfactual rule for assigning occupations, we need to predict each person's potential outcome under each counterfactual occupation they might have held under this rule. Doing so requires identifying the causal effect of occupation on the outcome of interest, for which we assume the causal directed acyclic graph (DAG) in Figure 5. Under this DAG, the health outcome person i would realize if counterfactually assigned to occupation d can be estimated by the predicted outcomes of people who are identical to person i along the measured covariates but who held the counterfactual occupation d. The credibility of causal identification is strengthened by the inclusion of relevant confounders: lagged health, demographic variables, and human capital variables.

There are two main threats to identification, both of which involve the coarse measurement of lagged health. The first threat is that self-reported lagged health

A) Make causal assumptions in a Directed Acyclic Graph (Pearl 2009)

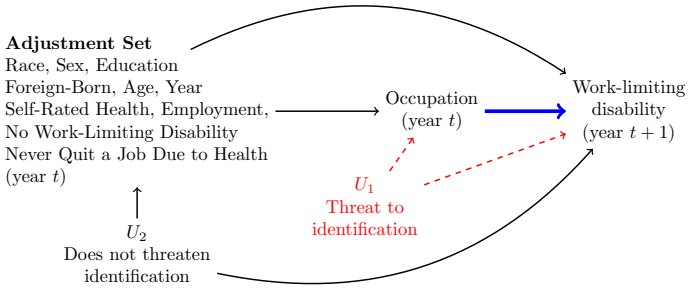

B) Fit a prediction function for the outcome

$$f(x, \vec{\ell}, d) = \mathrm{P}\left(Y = 1 \mid \text{Category} = x,\ \text{Covariates} = \vec{\ell},\ \text{Occupation} = d\right)$$

Probability of disability onset given racial category, covariates, and occupation

C) Predict counterfactuals for each unit i using the estimated function $\hat{f}()$.
Counterfactuals are defined over occupation only; keep other predictors as observed.
Predict for each occupation d.

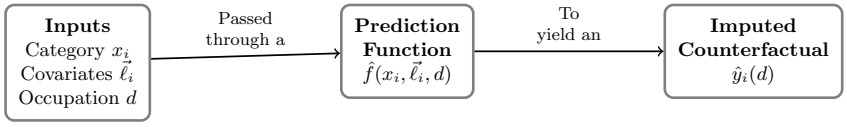

D) For each person i, average over the counterfactual occupation assignments

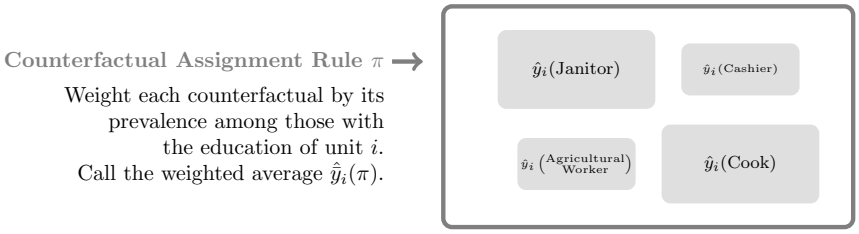

Example: Four occupations sized by prevalence among those who did not complete high school.

E) Report a sample-weighted average of $\hat{\bar{y}}_i(\pi)$ over people i in each racial category.

Figure 5: Procedure to estimate counterfactual disparities. Under causal assumptions (A), a predictive algorithm (B) can impute the counterfactual probability of work-limiting disability onset that each person i would realize if exposed to any occupation d (C). Weighting those counterfactuals by a new rule for assigning occupations (D) and averaging over subpopulations (E) yields counterfactual disparity estimates.

may be an insufficient control to block all confounding (online supplement Fig. 11). Health is a multi-dimensional and continuous construct, which may affect one's occupation and one's health outcomes through pathways not fully captured by a binary self-report of work-limiting disabilities. This threat exists in any observational study for which health is a confounder, and it is reduced but not eliminated by adjustment for several lagged measures of health: no work-limiting disability,

never quit a job due to health reasons, and five-point self-rated health. If we wish to study health in a causal framework without becoming fatalistic, a reasonable path forward is to adjust for many measures of health, as done in this article. The second threat to identification is that lagged health may be a pre-treatment collider (Greenland 2003; Greenland, Pearl, and Robins 1999, see the online supplement Fig. 12). Suppose that to be healthy at year t you either need to have held a safe occupation in the past or you need to have particularly robust health in unobserved ways. The less-robust people who held dangerous occupations in the past developed a work-limiting disability and were dropped out of the sample. Suppose also that the occupation you held in the past affects the occupation you hold in year t, due to persistence in one's occupation. In this case, among healthy people those in the most dangerous occupations would be those with the most robust unobserved health characteristics. This type of bias would make the study understate the degree to which dangerous occupations cause health problems, thus potentially understating the degree to which occupational segregation contributes to racial health disparities. Because lagged health is an important confounder, the risk of inducing this type of collider-stratification bias is likely to be outweighed by the bias-reducing benefits that arise from adjusting for lagged health. Overall, the two threats discussed here do not undermine the central approach taken in this article: by adjusting for measured confounders, the analyses marshal existing data to produce the best possible estimate of the causal contribution of occupational segregation to racial health disparities.

Estimate with Statistical Models

Under the causal assumptions, we can predict counterfactual outcomes using a model for the health outcome as a function of the adjustment set and occupation, as in past research on counterfactual disparities (Jackson and VanderWeele 2018; Lundberg 2024; VanderWeele and Robinson 2014). Because the analysis focuses on racial disparities, I allow all parameters to vary by race (with an r superscript) by estimating separate models by race. Because the treatment (occupation) can take hundreds of possible values, I estimate by a model with occupation random intercepts,

$$\mathrm{P}(Y = 1 \mid \text{Occupation} = o, \vec{X} = \vec{x}, \text{Race} = r) = \alpha_o^r + \vec{x}'\vec{\beta}^r \tag{2}$$

$$\alpha_o^r \sim \mathcal{N}(\vec{z_o}'\vec{\eta}^r, \sigma^2_{\alpha^r}) \tag{3}$$

where $\vec{x}$ are the measured covariates at the individual level, which include the adjustment set (operationalized linearly) as well as an indicator for years 2014 and later because of a questionnaire design change at that point (online supplement Section C). Note that the model implicitly conditions on lagged reports of no work-limiting disability (a lagged dependent variable) and never quitting a job for health reasons by restricting the sample along these variables. The second line is an occupation-level model for the occupation random intercepts, in which $\vec{z}_o$ is a vector containing the proportion Hispanic, non-Hispanic black, and other (with white omitted) in occupation o. Thus, the model allows occupations to have randomly varying effects that are regularized toward a linear function of the racial

composition of the occupation. The model is estimated using the `bam` function in the `mgcv` package in R (Wood 2017), which is designed to estimate a random intercept model quickly in a large sample size.

With the estimated prediction function, I then impute the counterfactual probability of work-limiting disability onset $\hat{y}_i(d)$ for each person i in each occupation d (Fig. 5C). Then, I average over a counterfactual assignment rule for occupations, weighting each potential outcome $\hat{y}_i(d)$ by the prevalence of the occupation d among people with the education level of person i (Fig. 5E). Finally, I aggregate the person-specific estimates over people in each subpopulation defined by race and ethnicity, using survey weights (Fig. 5E).

Interpret as an Intervention on Individuals, Not Societies

Because the analysis focuses on individuals, interpretations should likewise focus on individuals. We cannot say what would happen if the United States suddenly experienced a massive desegregation of occupations. Rather, we can draw a conclusion about individuals: if a few individuals were counterfactually exposed to a desegregated rule for assigning occupations, to what degree would a racial health disparity narrow? In short, the evidence collected at the level of individuals speaks to the causal role that macro-level segregation plays in the lives of individuals, leaving unanswered questions about how structures would shift under a massive reorganization of the labor market.

To draw broader societal-level conclusions from individual data would require arguments about the degree to which the intervention might change the mapping between treatments and outcomes (Jackson 2021), which might lead to encoding in a formal model (Jackson and Arah 2020). These are important questions for future research at the intersection of empirical evidence and theory. In the absence of strong theory about how systems would change if occupations desegregated, we focus here on the individual-level consequences of segregation on the health outcomes of individuals.

Results

Among the U.S. civilian population aged 25−60 in 2005−2020, work-limiting disabilities are common and disparate across categories of race/ethnicity: a work-limiting disability is reported for 13.1 percent of those who identify as non-Hispanic black but only 8.4 percent of non-Hispanic whites, 6.2 percent of Hispanics, and 6.1 percent of those who identify with another racial category (Fig. 6 Panel A). Focusing on the subpopulation who are employed and do not report a disability in year t, there are also disparities in the onset of disability by year $t + 1$, which is reported by 2.8 percent of non-Hispanic black workers compared with only 2.2 percent of non-Hispanic white workers, 2.0 percent of Hispanic workers, and 1.8 percent of those who identify with another racial category (Fig. 6 Panel B). Although the prevalence of disability onset (Panel B) is lower than the overall prevalence of disability (Panel A), focusing on onset is useful because it directs attention to the process by which those overall disparities arise.

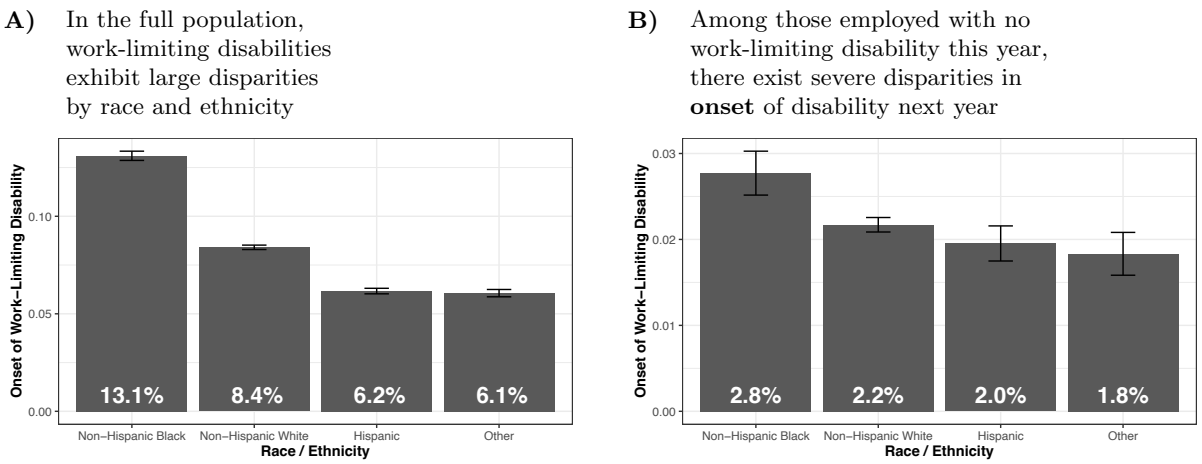

Figure 6: Disparities in work-limiting disability across population subgroups. Each panel pools data from the 2005 to 2020 March Supplement of the Current Population Survey. Error bars depict 95 percent confidence intervals based on variance estimates calculated using replicate weights (online supplement Section D).

To what degree are the disparities in Figure 6 Panel B a consequence of occupational segregation? Figure 7 summarizes the onset of work-limiting disability in the factual population and under the counterfactual assignment rule. Equalizing occupational assignments within education categories would reduce the onset of work-limiting disability among non-Hispanic black individuals from 2.7 percent to 2.6 percent and increase it slightly among non-Hispanic white individuals (who are now exposed to riskier occupations) (Panel A). The counterfactual black−white disparity is 61 percent smaller (95 percent CI: 27 percent, 56 percent) than the factual disparity: a reduction from a disparity of 0.59 percentage points to 0.35 percentage points (Panel B). A reduction of 0.24 percentage points corresponds to a large number of people: it implies that approximately 47,040 black workers who experienced onset of a work-limiting disability in a given year would not have done so if occupations were allocated equitably.[4] Unequal allocation to occupations is therefore responsible for some of the black−white disparity in the onset of work-limiting disability. It is noteworthy that such a large portion is accounted for even though the counterfactual assignment rule does not fully desegregate occupations but only desegregates them within educational categories. It is also noteworthy that the amount accounted for is not nearly the whole initial disparity. Structural racism shapes health in many forms, and other contributors to the racial disparity (e.g., neighborhood environments) expose people to differential hazards outside of work. Occupational segregation is a substantial part of the story, but it is not the whole story.

The two other counterfactual disparities (Hispanic vs. non-Hispanic white and other vs. non-Hispanic white) follow a different pattern. Each of these factual disparities runs in the opposite direction of the black−white disparity: these groups have *lower* onset of work-limiting disability compared to non-Hispanic whites. Yet they also work in more dangerous occupations. Equalizing occupational assignments would *increase* the Hispanic-white gap by 77 percent (95 percent CI: –10 percent, 263 percent) and the other-white gap by 42 percent (95 percent CI: 11 percent, 72 per-

A) Equalizing occupational assignments for a sample would reduce the onset of work-limiting disability for all groups except non-Hispanic whites

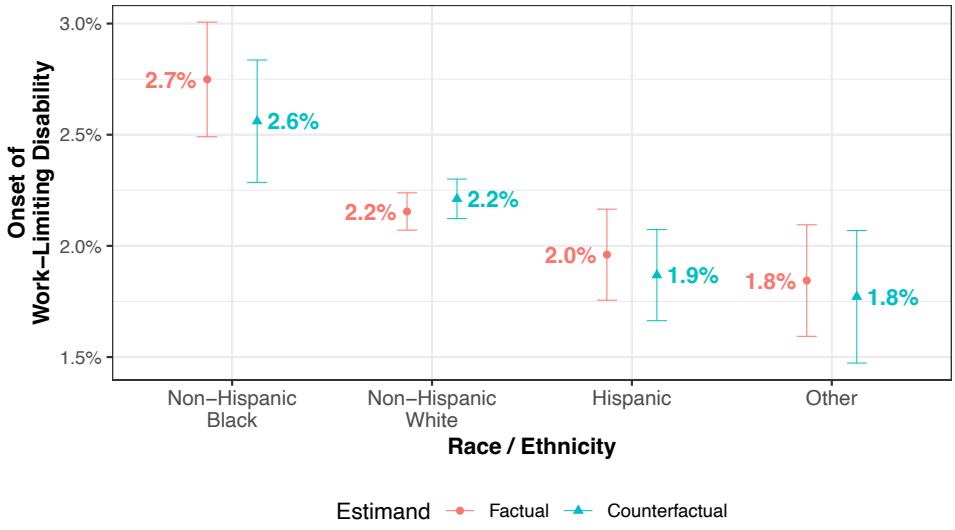

B) Equalizing occupational assignments for a sample would reduce the Black-white disparity by 61%. It would increase the health advantage of the other categories compared to non-Hispanic whites.

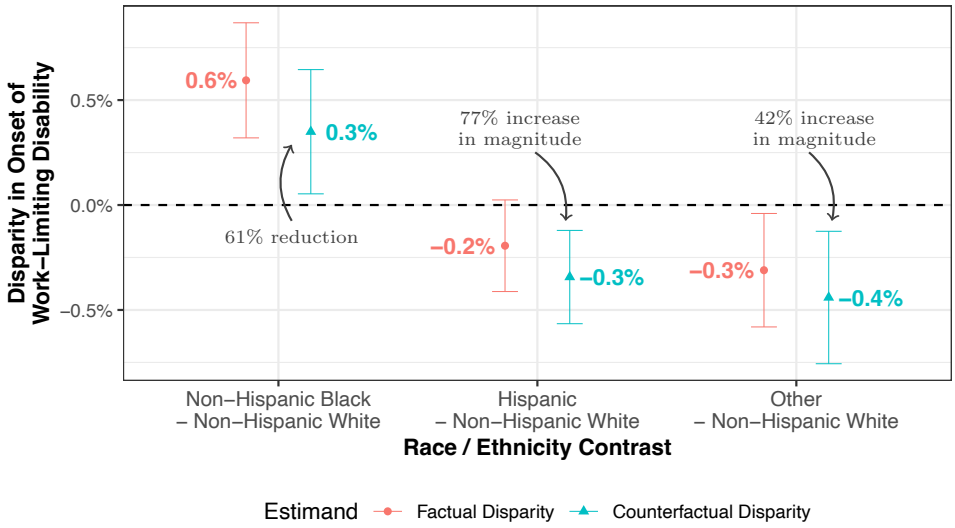

Figure 7: Factual and counterfactual probabilities of work-limiting disability onset. The counterfactual quantities refer to the subgroup average of the individual average potential outcome realized if occupations were assigned proportional to their prevalence within one's educational category. In Panel B, the black−white disparity reduction and other-white disparity increase noted with arrows are all statistically significantly different from zero at the 0.05 level. The Hispanic-white disparity increase is not statistically significant.

cent). The ability to make those claims is a key benefit of a causal perspective: even though these groups have descriptively better health than non-Hispanic whites, we would expect their health to be *even better* in a counterfactual world where occupations were allocated equitably. Hispanic individuals in the United States, for example, possess other covariates that are associated with good health, such as a younger age distribution and a higher proportion immigrants.[5] They do not fully realize the better health we might expect given those covariates in part because they are employed in more hazardous occupations.

Scope Conditions

Having defined and illustrated the method, this section revisits several scope conditions that future researchers should consider when applying this approach. The analysis requires the researcher to make choices that are not driven by data, including the definition of the causal question and the causal assumptions to answer that question. When researchers make these choices, they should convey to readers the resulting scope conditions on inferences.

First, the approach estimates the causal effect of one particular definition of segregation. The researcher operationalizes segregation when defining the treatment variable and potential outcomes. In the empirical illustration, segregation is across three-digit Census occupation codes. A counterfactual world that desegregates three-digit Census occupation codes does not eliminate segregation in other forms. Segregation may remain within the treatment categories: perhaps black taxi drivers serve areas that are more dangerous than the areas served by white taxi drivers. One could imagine a series of increasingly fine-grained social positions, with the number of treatments growing toward infinity. It is possible (though not guaranteed) that the effect of segregation would be larger if segregation were defined on a more fine-grained level with more treatment categories (e.g., occupations interacted with workplaces instead of just occupations). Researchers can only draw conclusions about the effect of segregation over the set of social positions that they define for the analysis. Researchers should also recognize that desegregation of one domain does not entail desegregation of other domains: people whose occupations are desegregated may still experience racial residential segregation, for instance, and their children may attend segregated schools. A benefit of the proposed approach is its transparency on this front: the researcher defines what they mean by segregation when they define the treatment and potential outcomes. Results only apply to this particular operationalization of segregation, and researchers should scope interpretations to the particular definition of segregation that they study.

Second, the approach is scoped to empirically tractable counterfactuals. Although we might like to know what would happen if occupations were fully desegregated, some counterfactual occupations are not empirically observed within confounder subgroups. People who did not graduate high school are never employed as high school teachers. As in research in social mobility that studies occupation as a causal treatment (Lundberg et al. 2024), data cannot answer questions about occupations that never happen within an educational group. The illustration avoided this problem by studying a world that desegregates occupations within

educational categories. Viewed as a limitation, this choice may bake in a substantial amount of occupational segregation that arises due to racial inequality in education. Viewed as an opportunity, this scope condition focuses the analysis on a particular step in the process that produces racial inequality. Unequal educational attainment is one step of the stratification process, and unequal occupational attainment within educational categories is a second step. Research on the first step is left as a separate topic of research, for which the causal exposure variable should be education rather than occupation. When researchers develop a counterfactual assignment rule that is constrained to empirically tractable counterfactuals, they should be transparent about the sources of the disparity that are not addressed by their analysis. In this case, the focus is on occupational segregation within educational categories rather than occupational segregation as a whole.

Third, the approach assumes that the prediction function mapping predictors to potential outcomes is the same in the factual (segregated) setting as in the counterfactual (desegregated) setting. As an example in which that would fail, consider a counterfactual in which the number of white roofers rises dramatically. White people might then use their political power to enact new regulations to improve the safety of roofing, for example, requiring improved safety harnesses. The prediction function mapping roofing to health risks would change. In other words, people with resources who are exposed to a new health risk may use their resources to reduce the danger associated with that risk, which is one reason social conditions have long been theorized as fundamental causes of disease (Link and Phelan 1995). General equilibrium dynamics where the mapping changes when many people's treatments change are an important area for future research (Jackson and Arah 2020). An implication for the present approach is that researchers should be transparent that inference is about counterfactuals defined over individuals, not societies, as discussed earlier in the article. If an individual's social position changes, that individual is not powerful enough to change social structures. Without studying what would happen if everyone's social positions changed all at once, one can still study a more local claim about how the existing structure shapes the outcomes of individuals on average.

Fourth, inference is only as credible as the causal identification assumptions. As with all observational causal inference, results could be misleading if individuals select into the social positions (e.g., occupations) as a function of their potential outcomes (e.g., their potential health in a particular occupation). The prediction function learned in the factual world would then not apply in the counterfactual world. An explicitly causal framework makes this limitation transparent, because the researcher must state and defend causal assumptions, for example, with a DAG. When interpreting results, researchers should bear in mind that inference is only as credible as those assumptions.

Discussion

It is conceptually plausible that social disparities may be caused in part by segregation, whether across occupations, schools, or neighborhoods. Yet studies rarely carry out analyses to explicitly consider the causal effects of segregation. One reason

causal studies are uncommon may be that segregation is often conceptualized as a characteristic of society, whereas many data sets contain variation across individuals within a single society. This article showed how to study occupational segregation as a cause of disparities using potential outcomes defined at the individual level.

The approach of this article grew from a concrete substantive question: to what degree does racial occupational segregation cause racial disparities in reporting a work-limiting disability? By analyzing panel data under a set of causal assumptions, results showed that the black−white gap in the onset of work-limiting disability would close by 61 percent in a counterfactual world where individuals were exposed to occupations by a rule that depended only on their education. Substantively, this result emphasizes the importance of occupational segregation for understanding racial disparities in health.

The primary contribution of this article is methodological. Although sociologists have begun to study the causal inputs that shape disparities, past work has focused on binary treatment variables (e.g., Lundberg 2024). The present article broadens those ideas to study how disparities can be caused by segregation across a categorical exposure (occupation). Defining the causal effect of occupational segregation in the potential outcomes framework is especially useful because it bridges questions in sociology often posed at the macro level (segregation) to methodological approaches in the potential outcomes framework that are more often applied to purely micro-level questions. There will always be an important place for macro-level analyses of macro-level phenomena. In addition to that well-studied path, the approach of this article shows how microdata analyzed in a micro-level causal framework can answer questions relevant to macro-level claims. One can study how individuals respond causally to changes in the individual-level exposure (occupation) that collectively comprises a macro-level phenomenon (segregation). Future studies will hopefully build on this example to answer causal questions with microdata that speak to causal effects that may help us understand macro-level causes such as segregation.

Notes

1 This measurement of work-limiting disability is widely used in studies using the CPS (e.g., Brucker et al. 2015; Jajtner et al. 2020) as well as studies relying on similar questions in the Health and Retirement Study (e.g., Burkhauser et al. 1999), the Panel Study of Income Dynamics (e.g., Meyer and Mok 2019), and the National Longitudinal Survey of Youth (e.g., Pais 2014). The online supplement Section C discusses limitations of this operationalization of work-limiting disability: it only captures respondent perceptions of their own health, and the CPS survey instrument changed slightly over time.

2 Supplementary analyses contain a more thorough adjustment set at the cost of narrowing the scope of years involved. Beginning in 2009, the CPS added a series of six questions in which respondents reported whether they had any difficulty with hearing, vision, remembering, walking, or climbing stairs, performing basic activities outside the home alone, or taking care of their personal needs. Analyses on the subset of people who answer no to all of those questions yield substantively similar results (online supplement Fig. 16).

3 The empirical analysis focuses on the 428 occupations for which the data contain people of all four racial categories studied.

4 In the fourth quarter of 2019, there were 19.6 million employed black Americans (Edwards and Smith 2020). The above estimate (based on 2005−2020) implies that 0.24 percent of these individuals experience disability onset and would not have experienced that onset if occupations were allocated equitably. That corresponds to 0.24 percent × 19.6 million = 47,040 people.

5 In an alternative specification restricted to the native born population, the onset of work-limiting disabilities is factually slightly higher for Hispanics compared with non-Hispanic whites (online supplement Fig. 15).

References

Ahonen, Emily Q., Kaori Fujishiro, Thomas Cunningham, and Michael Flynn. 2018. "Work as an Inclusive Part of Population Health Inequities Research and Prevention". *American Journal of Public Health* 108(3):306–11. https://doi.org/10.2105/AJPH.2017.304214

Bearman, Peter S., James Moody, and Katherine Stovel. 2004. "Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks". *American Journal of Sociology* 110(1):44–91. https://doi.org/10.1086/386272

Blau, Francine D. and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations". *Journal of Economic Literature* 55(3):789–865. https://doi.org/10.1257/jel.20160995

Brucker, Debra L., Sophie Mitra, Navena Chaitoo, and Joseph Mauro. 2015. "More Likely to Be Poor Whatever the Measure: Working-Age Persons with Disabilities in the United States". *Social Science Quarterly* 96(1):273–96. https://doi.org/10.1111/ssqu.12098

Buchanan, Susan, Pamela Vossenas, Niklas Krause, Joan Moriarty, Eric Frumin, Jo Anna M. Shimek, Franklin Mirer, Peter Orris, and Laura Punnett. 2010. "Occupational Injury Disparities in the U.S. Hotel Industry". *American Journal of Industrial Medicine* 53(2):116–25. https://doi.org/10.1002/ajim.20724

Bureau of Labor Statistics. 2020. "Employer-Reported Workplace Injuries and Illnesses–2019". *U.S. Department of Labor News Release*, 20–2030.

Burgard, Sarah A. and Katherine Y. Lin. 2013. "Bad Jobs, Bad Health? How Work and Working Conditions Contribute to Health Disparities". *American Behavioral Scientist* 57(8):1105–27. https://doi.org/10.1177/0002764213487347

Burkhauser, Richard V., J.S. Butler, Yang-Woo Kim, and Robert R. Weathers. 1999. "The Importance of Accommodation on the Timing of Disability Insurance Applications: Results from the Survey of Disability and Work and the Health and Retirement Study". *Journal of Human Resources* 589–611. https://doi.org/10.2307/146381

Coleman, James S. 1990. *Foundations of Social Theory*. Harvard University Press.

Edwards, Roxanna and Sean M. Smith. 2020. "Job Market Remains Tight in 2019, as the Unemployment Rate Falls to Its Lowest Level Since 1969". *Bureau of Labor Statistics: Monthly Labor Review*, April. https://doi.org/10.21916/mlr.2020.8

Ferguson, John-Paul and Rembrand Koning. 2018. "Firm Turnover and the Return of Racial Establishment Segregation". *American Sociological Review* 83(3):445–74. https://doi.org/10.1177/0003122418767438

Fleischer, Nancy L., Hope M. Tiesman, Jeri Sumitani, Terry Mize, Kumar Kartik Amarnath, A. Rana Bayakly, and Matthew W. Murphy. 2013. "Public Health Impact of Heat-Related Illness among Migrant Farmworkers". *American Journal of Preventive Medicine* 44(3):199–206. https://doi.org/10.1016/j.amepre.2012.10.020

Flippen, Chenoa. 2004. "Unequal Returns to Housing Investments? A Study of Real Housing Appreciation among Black, White, and Hispanic Households". *Social Forces* 82(4):1523–51. https://doi.org/10.1353/sof.2004.0069

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2020. *Integrated Public Use Microdata Series, Current Population Survey: Version 8.0 [Dataset].* Minneapolis, MN: IPUMS.

Greenland, Sander. 2003. "Quantifying Biases in Causal Models: Classical Confounding Vs Collider-Stratification Bias". *Epidemiology* 14(3):300–6. https://doi.org/10.1097/01.EDE.0000042804.12056.6C

Greenland, Sander, Judea Pearl, and James M. Robins. 1999. "Causal Diagrams for Epidemiologic Research". *Epidemiology* 10(1):37–48. https://doi.org/10.1097/00001648-199901000-00008

Harknett, Kristen, Daniel Schneider, and Rebecca Wolfe. 2020. "Losing Sleep over Work Scheduling? The Relationship between Work Schedules and Sleep Quality for Service Sector Workers". *SSM-Population Health* 12:100681. https://doi.org/10.1016/j.ssmph.2020.100681

Hedström, Peter and Peter S. Bearman. 2009. *The Oxford Handbook of Analytical Sociology*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199215362.001.0001

Heymann, Jody. 2003. *Global Inequalities at Work: Work's Impact on the Health of Individuals, Families, and Societies*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195150865.001.0001

Imbens, Guido W. and Donald B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press. https://doi.org/10.1017/CBO9781139025751

Jackson, John W. 2021. "Meaningful Causal Decompositions in Health Equity Research: Definition, Identification, and Estimation Through a Weighting Framework". *Epidemiology* 32(2):282–90. https://doi.org/10.1097/EDE.0000000000001319

Jackson, John W. and Onyebuchi A. Arah. 2020. "Invited Commentary: Making Causal Inference More Social and (Social) Epidemiology More Causal". *American Journal of Epidemiology* 189(3):179–82. https://doi.org/10.1093/aje/kwz199

Jackson, John W. and Tyler J. VanderWeele. 2018. "Decomposition Analysis to Identify Intervention Targets for Reducing Disparities". *Epidemiology* 29(6):825–35. https://doi.org/10.1097/EDE.0000000000000901

Jajtner, Katie M., Sophie Mitra, Christine Fountain, and Austin Nichols. 2020. "Rising Income Inequality Through a Disability Lens: Trends in the United States 1981–2018". *Social Indicators Research* 151(1):81–114. https://doi.org/10.1007/s11205-020-02379-8

Juon, Hee-Soon, Alicia Hong, Marcella Pimpinelli, Madhuwani Rojulpote, Russell McIntire, and Julie A. Barta. 2021. "Racial Disparities in Occupational Risks and Lung Cancer Incidence: Analysis of the National Lung Screening Trial". *Preventive Medicine* 143:106355. https://doi.org/10.1016/j.ypmed.2020.106355

Link, Bruce G. and Jo Phelan. 1995. "Social Conditions as Fundamental Causes of Disease". *Journal of Health and Social Behavior* 80–94. https://doi.org/10.2307/2626958

Lipscomb, Hester J., Dana Loomis, Mary Anne McDonald, Robin A. Argue, and Steve Wing. 2006. "A Conceptual Model of Work and Health Disparities in the United States". *International Journal of Health Services* 36(1):25–50. https://doi.org/10.2190/BRED-NRJ7-3LV7-2QCG

Lowry, Sarah J., Hillary Blecker, Janice Camp, Butch De Castro, Steven Hecker, Saman Arbabi, Neal Traven, and Noah S. Seixas. 2010. "Possibilities and Challenges in Occupational Injury Surveillance of Day Laborers". *American Journal of Industrial Medicine* 53(2):126–34. https://doi.org/10.1002/ajim.20741

Lundberg, Ian. 2024. "The Gap-Closing Estimand: A Causal Approach to Study Interventions That Close Disparities across Social Categories". *Sociological Methods & Research* 53(2):507–70. https://doi.org/10.1177/00491241211055769

Lundberg, Ian, Daniel Molitor, and Jennie E. Brand. 2025. "The Causal Effect of Parent Occupation on Child Occupation: A Multivalued Treatment with Positivity Constraints." *Sociological Methods & Research* 54(4):1435–1462. https://doi.org/10.31235/osf.io/rkw3e

Markley, Scott N., Taylor J. Hafley, Coleman A. Allums, Steven R. Holloway, and Hee Cheol Chung. 2020. "The Limits of Homeownership: Racial Capitalism, Black Wealth, and the Appreciation Gap in Atlanta". *International Journal of Urban and Regional Research* 44(2):310–28. https://doi.org/10.1111/1468-2427.12873

Meyer, Bruce D. and Wallace K.C. Mok. 2019. "Disability, Earnings, Income and Consumption". *Journal of Public Economics* 171:51–69. https://doi.org/10.1016/j.jpubeco.2018.06.011

Moen, Phyllis, Erin L. Kelly, Eric Tranby, and Qinlei Huang. 2011. "Changing Work, Changing Health: Can Real Work-Time Flexibility Promote Health Behaviors and Well-Being?" *Journal of Health and Social Behavior* 52(4):404–29. https://doi.org/10.1177/0022146511418979

Ólafsdóttir, Sigrun and Jason Beckfield. 2020. "Health Inequalities: What Do We Know and What Do We Need to Know?" *Social Science & Medicine* 267:113575. https://doi.org/10.1016/j.socscimed.2020.113575

Pais, Jeremy. 2014. "Cumulative Structural Disadvantage and Racial Health Disparities: The Pathways of Childhood Socioeconomic Influence". *Demography* 51(5):1729–53. https://doi.org/10.1007/s13524-014-0330-9

Pearl, Judea. 2009. *Causality*. Cambridge University Press.

Petersen, Trond and Laurie A. Morgan. 1995. "Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap". *American Journal of Sociology* 101(2):329–65. https://doi.org/10.1086/230727

Prins, Seth J., Sarah McKetta, Jonathan Platt, Carles Muntaner, Katherine M. Keyes, and Lisa M. Bates. 2019. "Mental Illness, Drinking, and the Social Division and Structure of Labor in the United States: 2003-2015". *American Journal of Industrial Medicine* 62(2):131–44. https://doi.org/10.1002/ajim.22935

Rivera Drew, Julia A., Sarah Flood, and John Robert Warren. 2014. "Making Full Use of the Longitudinal Design of the Current Population Survey: Methods for Linking Records Across 16 Months". *Journal of Economic and Social Measurement* 39(3):121–44. https://doi.org/10.3233/JEM-140388

Sampson, Robert J. and Alix S. Winter. 2016. "The Racial Ecology of Lead Poisoning: Toxic Inequality in Chicago Neighborhoods, 1995-2013". *Du Bois Review: Social Science Research on Race* 13(2):261–83. https://doi.org/10.1017/S1742058X16000151

Schulte, Paul A., Rene Pana-Cryan, Teresa Schnorr, Anita L. Schill, Rebecca Guerin, Sarah Felknor, and Gregory R. Wagner. 2017. "An Approach to Assess the Burden of Work-Related Injury, Disease, and Distress". *American Journal of Public Health* 107(7):1051–57. https://doi.org/10.2105/AJPH.2017.303765

Seabury, Seth A., Sophie Terp, and Leslie I. Boden. 2017. "Racial and Ethnic Differences in the Frequency of Workplace Injuries and Prevalence of Work-Related Disability". *Health Affairs* 36(2):266–73. https://doi.org/10.1377/hlthaff.2016.1185

Stainback, Kevin and Donald Tomaskovic-Devey. 2012. *Documenting Desegregation: Racial and Gender Segregation in Private Sector Employment Since the Civil Rights Act*. Russell Sage Foundation.

Stanbury, Martha and Kenneth D. Rosenman. 2014. "Occupational Health Disparities: A State Public Health-Based Approach". *American Journal of Industrial Medicine* 57(5):596–604. https://doi.org/10.1002/ajim.22292

Tomaskovic-Devey, Donald, Catherine Zimmer, Kevin Stainback, Corre Robinson, Tiffany Taylor, and Tricia McTague. 2006. "Documenting Desegregation: Segregation in American Workplaces by Race, Ethnicity, and Sex, 1966–2003". *American Sociological Review* 71(4):565–88. https://doi.org/10.1177/000312240607100403

Torrats-Espinosa, Gerard. 2021. "Using Machine Learning to Estimate the Effect of Racial Segregation on COVID-19 Mortality in the United States". *Proceedings of the National Academy of Sciences* 118(7):e2015577118. https://doi.org/10.1073/pnas.2015577118

VanderWeele, Tyler J. and Whitney R. Robinson. 2014. "On Causal Interpretation of Race in Regressions Adjusting for Confounding and Mediating Variables". *Epidemiology* 25(4):473–84. https://doi.org/10.1097/EDE.0000000000000105

Weathers, Ericka S. and Victoria E. Sosina. 2022. "Separate Remains Unequal: Contemporary Segregation and Racial Disparities in School District Revenue". *American Educational Research Journal* 59(5):905–38. https://doi.org/10.3102/00028312221079297

Wood, Simon N. 2017. *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC Press. https://doi.org/10.1201/9781315370279

Acknowledgments: For helpful discussions and feedback relevant to this project, I thank Brandon Stewart, Matthew Salganik, Dalton Conley, Sara McLanahan, Rebecca Johnson, Gillian Slee, and participants in presentations at the Princeton Department of Sociology, the Cornell Center for the Study of Inequality, and the UCSF Department of Social and Behavioral Sciences, as well as the editor and anonymous reviewers. The author benefited from facilities and resources provided by the California Center for Population Research at UCLA (CCPR), which receives core support (P2C-HD041022) from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD). The content is solely the responsibility of the authors and does not necessarily represent the official views of the Eunice Kennedy Shriver National Institute of Child Health & Human Development or the National Institutes of Health.

Ian Lundberg: Department of Sociology, UCLA.
E-mail: ianlundberg.org, ianlundberg@ucla.edu.