Intersectional Privilege and Aggregate Wage Inequality

Jan Schulz[1], Caleb E. Agoha[2], and Daniel M. Mayerhoffer[3]

[1]Economics Institute, University of Bamberg, Bamberg, Germany
[2]Oxford Internet Institute, University of Oxford, Oxford, United Kingdom
[3]Amsterdam Institute for Social Science Research and Data Science Centre, University of Amsterdam, Amsterdam, The Netherlands

Corresponding author: Jan Schulz, jan.schulz-gebhard@uni-bamberg.de

Abstract

Research on intersectional wage gaps usually asks whether workers who combine marginalised identities face a wage penalty beyond the separate penalties associated with each identity. We examine the same nonadditivity from the opposite corner of the wage structure and ask how it relates to aggregate wage inequality. Using 26.5 million observations from the 1980–2024 U.S. Census and American Community Survey, we construct a no-interaction wage benchmark for White men from the observed wages of White women, Black or Hispanic men, and Black or Hispanic women. The resulting White-male premium is positive over almost the full wage distribution, particularly at higher wage ranks, and remains economically significant after conditioning on worker and job characteristics. Removing this interaction lowers the Gini coefficient of the complete eligible-worker distribution by about 3 per cent. Redistributing the excess privilege to the other groups leads to substantially larger reductions. The results connect intersectional wage structure to aggregate inequality and highlight the distributional role of multiply advantaged groups.

Keywords: intersectionality; wage inequality; racial wage gaps; gender wage gaps; privilege; Gini coefficient

JEL Classification: D31; J15; J16; J31

1 Introduction

Gender and racial wage gaps are usually studied one dimension at a time. Intersectional research starts from the observation that this can be misleading: the wage position of Black women, for example, need not be recovered by adding a "gender effect" to a "race effect." The empirical literature has therefore examined whether workers who combine marginalised identities face an additional wage penalty beyond the separate disadvantages associated with each identity. Less attention has been paid to what such nonadditivity implies for the overall distribution of wages. Yet an interaction between race and sex changes not only relative wages across groups but also the location of wage mass in the aggregate distribution. Its consequences for inequality therefore depend on which group receives the nonadditive component and where that group is located in the wage distribution.

The same four-cell wage structure used to study cumulative disadvantage can also be represented from the perspective of the multiply advantaged group. If the proportional gender gap is larger among White workers than among Black or Hispanic workers, or equivalently if the racial wage advantage is larger among men than among women, the interaction can be allocated to White men as an excess wage premium relative to a no-interaction benchmark. This representation is particularly useful for studying aggregate inequality because it identifies a specific component of the wage structure that can be removed while leaving the remaining group wage profiles unchanged.

This paper asks how much aggregate wage inequality is associated with this nonadditive component of race-by-sex wage differences. Using IPUMS USA microdata from the 1980, 1990, and 2000 Censuses and annual American Community Survey samples from 2005 to 2024 (excluding 2020), we construct a no-interaction wage benchmark for White men from the observed wages of White women, minority[1] men, and minority women. We do so separately for Black/White and Hispanic/White comparisons and follow the resulting interaction across the wage distribution and over time. We then apply the same benchmark to individual wages and recompute inequality for the complete eligible-worker population, including workers outside the four groups used to identify the interaction.

The resulting link to aggregate inequality is substantial. The White-male premium is positive over almost the entire wage distribution and is particularly pronounced at higher wage ranks. Removing the interaction lowers the Gini coefficient of the complete worker distribution by about 3 per cent in 2024 in both comparisons. Because this counterfactual also reduces total wage mass, we additionally consider a redistribution exercise that transfers the removed wage amount to the other three groups while holding aggregate wage mass constant. The resulting decline in inequality is considerably larger. The central result is therefore not simply that a nonadditive wage premium can be identified, but that its size, its attachment to a large group, and its position in the upper part of the wage distribution make it quantitatively relevant for aggregate wage inequality.[2]

The pattern also survives adjustment for age, education, usual hours, state, metropolitan status, occupation, and industry. Residualisation reduces the estimated premium, but the interaction remains positive across the wage distribution and continues to generate a sizeable proportional change in the Gini. The aggregate inequality result is therefore not driven solely by observed differences in worker characteristics or by the occupational and industrial composition of employment.

The analysis connects two literatures that have rarely been brought together directly. Quantitative intersectionality has documented nonadditive wage differences across race, ethnicity, gender, class, and nativity (Almquist, 1975; Greenman and Xie, 2008; Kim, 2009; Mandel and

[1]We use "minority" as shorthand for Black or Hispanic workers in the U.S. context, depending on the comparison. This terminology neither encompasses all U.S. minority groups nor implies that numerical minority status itself explains the observed wage differences or any underlying discrimination.

[2]The redistribution exercise is an accounting counterfactual. It does not incorporate behavioural or general-equilibrium responses and should not be interpreted as an estimate of the effect of a feasible tax reform.

Semyonov, 2016; George et al., 2022; Paul et al., 2022; Sprengholz and Hamjediers, 2024). Recent work has also made privilege itself an explicit object of analysis (Cech, 2022; Alonso-Villar and del Río, 2023). The inequality literature, meanwhile, asks how differences across population groups map into aggregate inequality (Shorrocks, 1984; Cowell and Jenkins, 1995; Heikkuri and Schief, 2026). Our objective differs from a conventional subgroup decomposition. We do not ask what share of the Gini is "due to race" or "due to gender." We isolate the nonadditive component of a particular four-group wage structure, allocate it to the multiply advantaged cell, and trace its consequences through the complete wage distribution. In this way, the intersectional interaction becomes an object of inequality analysis rather than only a feature of a wage-gap regression.

The paper also extends earlier work on the reference-group problem in intersectional wage measurement (Schulz et al., 2026). That paper focuses on group means and on alternative penalty and privilege representations of the same four-cell wage structure. The present analysis uses individual microdata, follows the interaction across the wage distribution and over four decades, incorporates survey-design inference, and makes aggregate inequality the main outcome. It also establishes how the inequality effect depends on the location of the associated wage mass in the pooled distribution. The paper therefore links the measurement of intersectional nonadditivity to a broader question about the sources and structure of wage inequality.

2 Intersectional nonadditivity and the wage distribution

2.1 From cumulative disadvantage to relational wage structure

Intersectionality was introduced to describe social positions that cannot be understood by treating axes of stratification as independent (Crenshaw, 1989). Quantitative wage research has operationalised this insight in several ways, including interaction terms, pairwise decompositions, matching, reweighting, and comparisons across detailed intersectional groups. An early question was whether Black women experience a "double jeopardy" that exceeds the separate disadvantages associated with race and gender (Almquist, 1975). Later work found that the answer depends on the comparison and specification. Greenman and Xie (2008), for example, report that minority women often face a smaller gender penalty than White women, whereas Kim (2009), George et al. (2022), and Paul et al. (2022) document settings in which Black women bear an additional penalty. Using a broader set of marginalised identities, McLaughlin and Neumark (2026) find little evidence that wage penalties are generally more than additive.

Long-run U.S. evidence also shows why gender cannot simply be treated as a fixed overlay on the racial pay gap. Using IPUMS data for 1970–2010, Mandel and Semyonov (2016) find that racial pay gaps are larger among men and that, after decades of convergence, both gross and adjusted gaps widened again after 2000. Their object is the racial pay gap estimated separately by gender. Ours is the difference between those gender-specific gaps. The two are related, but they need not move in the same direction.

The apparent disagreement is partly substantive: different groups, years, outcomes, samples, and controls need not generate the same interaction. It is also partly a matter of what an interaction is asked to represent. A four-cell interaction is relational. When the interaction is positive in the log-wage contrast used below, the wage difference between White men and White women exceeds the corresponding male-female difference among minority workers. The same fact can be described as a smaller gender penalty among minority workers, a larger racial advantage among men, or an excess premium in the White-male cell. None of these descriptions, by itself, identifies the causal mechanism that produced the observed wages.

Our privilege-centred reading has precedents. Alonso-Villar and del Río (2023) separate conditional earnings advantages and disadvantages across twelve U.S. gender-by-race/ethnicity groups, while Cech (2022) documents a broad intersectional advantage for White, able-bodied, heterosexual men in STEM. In a four-group German setting, Sprengholz and Hamjediers (2024) find that much of the unexplained wage difference is concentrated in comparisons involving native men and explicitly describe the pattern as a double advantage rather than a double disadvantage. Other recent work broadens the set of intersectional positions used to study income and wage inequality across race, gender, class, nativity, ethnicity, and geography (Vo et al., 2023; Restifo et al., 2023; Alonso-Villar and del Río, 2024; Reiss et al., 2025; Madsen et al., 2026). These studies emphasise heterogeneity across intersectional groups. Our question is different in its aggregation: we isolate one nonadditive component and ask how it changes a standard inequality index for the complete wage distribution. Our use of "privilege" is correspondingly narrow. It refers to the wage amount allocated to the multiply advantaged cell by a transparent no-interaction accounting benchmark.

Thus, the empirical premium can reflect discrimination, bargaining institutions, occupational segregation, work organisation, selection, unmeasured productivity, or other mechanisms. Existing research points to several potential sources of race-by-gender wage differences. Educational attainment accounts for only part of these disparities (Budig et al., 2021), while occupational segregation and social closure contribute to both gender and racial wage gaps (Tomaskovic-Devey, 1993). Nonlinear returns to long and inflexible hours may also help explain why gender wage gaps are particularly pronounced at the upper end of the wage distribution (Goldin, 2014). Longitudinal employer-employee data further show that differences in returns to characteristics become especially important in the upper part of the earnings distribution (McKinney et al., 2022). Our design does not attempt to adjudicate among these explanations but asks what the observed nonadditivity contributes to the distribution in which it appears.

2.2 The inequality question

The usual subgroup-decomposition problem starts with an inequality index and partitions it into within-group and between-group components. For additively decomposable measures, this has a long axiomatic foundation (Shorrocks, 1984). Related work asks how much of observed inequality can be explained by population characteristics (Cowell and Jenkins, 1995). For the Gini coefficient in particular, overlap among subgroup distributions has made exact decomposition difficult. Heikkuri and Schief (2026) provide a recent axiomatic within/between decomposition that resolves this problem.

Our exercise is complementary. A within/between decomposition describes how much aggregate inequality is associated with differences across the full set of population subgroups. It does not isolate the interaction between two group dimensions inside a four-cell table. Conversely, the race-by-sex interaction is not itself an inequality index but a local feature of group wage profiles. To connect the two, we use the interaction to define a counterfactual wage schedule, apply it to individuals in the full population, and recompute the Gini.

Distributional approaches to wage discrimination have likewise stressed that regression-based mean gaps and distribution-wide measures answer different questions. del Río et al. (2011) discuss limitations of the Oaxaca–Blinder decomposition and of approaches based on quantile regressions or counterfactual functions, and propose measuring the distribution of individual wage gaps directly. More recently, De Poli and Maier (2025) use microsimulation to trace estimated gender pay penalties into labour-earnings inequality, household income inequality, and poverty. Our exercise shares the move from a group-specific wage structure to an aggregate distributional outcome, but its object is different. Rather than estimating a discrimination component, we take the observed race-by-sex interaction as the accounting object, allocate it to the White-male cell through the no-interaction benchmark, and evaluate the resulting change in the pooled wage distribution.

This distinction is useful because intersectional structure can matter even when one-dimensional group gaps are modest. Suppose, for example, that White men receive an additional premium only in the upper half of their wage distribution. A decomposition based solely on mean gender or racial gaps can miss where the associated wage mass sits in the pooled distribution. The Gini is particularly sensitive to this location: an additional dollar paid at a high pooled rank has a different inequality effect from an additional dollar paid near the centre. Relatedly, Giangregorio and Villani (2025) show that the inequality effect of a marginal wage increase depends on the distributional position of the group receiving it.

For our counterfactual, this dependence on location can be stated exactly. Let $t \in [0, 1]$ scale the premium from the no-interaction distribution to the observed distribution, let $G(t)$ denote the pooled Gini, and let $\overline{R}^D(t)$ denote the premium-weighted pooled rank at which the additional White-male wage mass is received. If the individual premiums are non-negative and their total is positive, then

$$\frac{\partial G(t)}{\partial t} > 0 \quad \Longleftrightarrow \quad \overline{R}^D(t) > \frac{1 + G(t)}{2}. \tag{1}$$

The premium therefore raises the Gini locally when its wage mass is concentrated sufficiently high in the pooled distribution. For example, if $G(t) = 0.4$, the threshold is the 70th percentile. Appendix F provides the derivation and shows how the local result aggregates to the difference between the observed and no-interaction Ginis.

2.3 The privilege benchmark

Let $c \in \{B, H\}$ denote the comparison branch. For the Black/White branch, $cM = BM$ and $cW = BW$; for the Hispanic/White branch, $cM = HM$ and $cW = HW$. Let Y^g denote a wage outcome for group g. It may be a mean wage, a wage at a common within-group percentile, or a rank-bin mean used in the aggregate counterfactual.

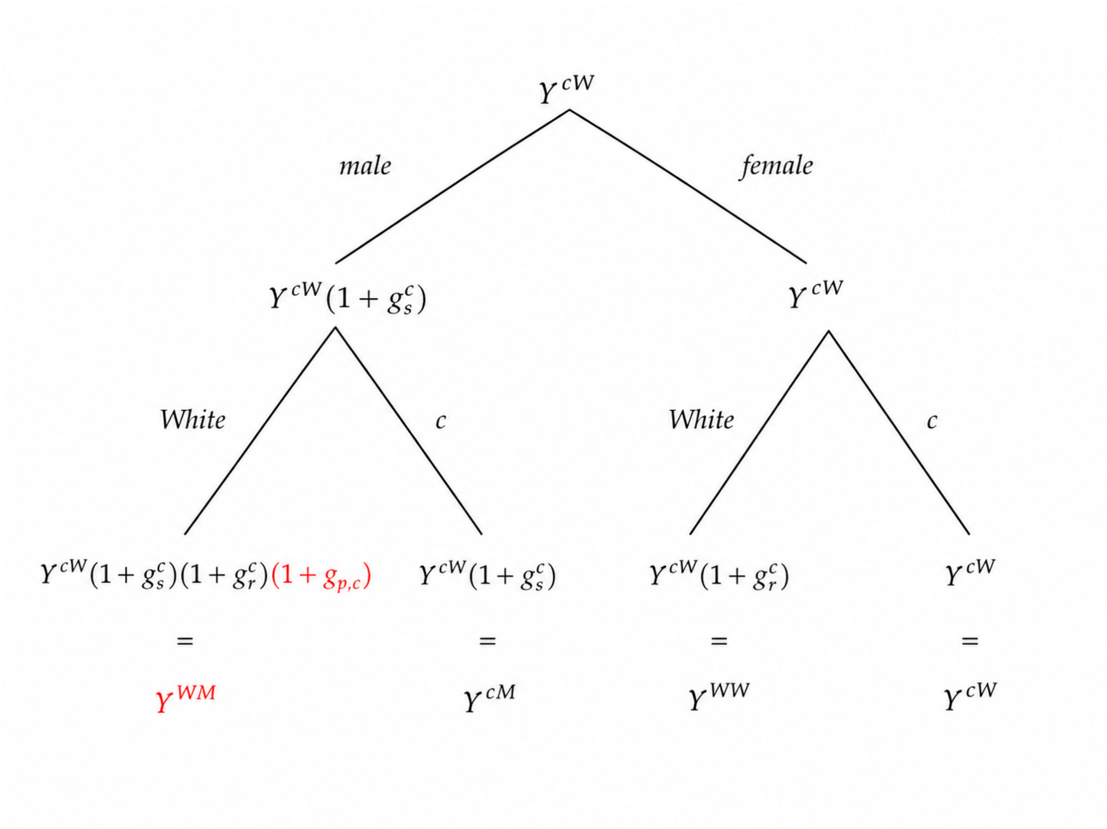

Fig. 1 Privilege tree for branch c. The sex and White/minority factors reproduce the other three group outcomes. The additional factor $1+g_{p,c}$ is the multiplicative interaction allocated to the White-male path

We start from minority women. The male/female wage factor among minority workers is

$$1+g_s^c = \frac{Y^{cM}}{Y^{cW}}, \tag{2}$$

and the White/minority factor among women is

$$1+g_r^c = \frac{Y^{WW}}{Y^{cW}}. \tag{3}$$

If these two factors were sufficient to reproduce White-male wages, the no-interaction outcome would be

$$\widehat{Y}_{0,c}^{WM} = Y^{cW}(1+g_s^c)(1+g_r^c) = \frac{Y^{cM}Y^{WW}}{Y^{cW}}. \tag{4}$$

We define the proportional excess White-male premium relative to this benchmark as

$$g_{p,c} = \frac{Y^{WM} - \widehat{Y}_{0,c}^{WM}}{\widehat{Y}_{0,c}^{WM}} = \frac{Y^{cW}Y^{WM}}{Y^{cM}Y^{WW}} - 1. \tag{5}$$

The same relationship is displayed in Figure 1. The interaction term appears only on the White-male path, while the other three observed outcomes determine the no-interaction benchmark.

Equation (5) is equivalent to a conventional two-way interaction in logarithms:

$$\log(1+g_{p,c}) = \log Y^{WM} - \log Y^{WW} - \log Y^{cM} + \log Y^{cW}. \tag{6}$$

The no-interaction condition is therefore $g_{p,c} = 0$. A positive value means that observed White-male wages exceed the level implied by the sex ratio among minority workers and the White/minority ratio among women. We use the benchmark in Equation (4) throughout; in particular, we do not normalise the premium by observed White-male wages.

For distributional comparisons, let $Q_g(p,t)$ be the weighted wage quantile of group g at within-group rank p in year t. The rank-indexed benchmark is

$$\widehat{Q}^{WM}_{0,c}(p,t) = \frac{Q_{cM}(p,t)Q_{WW}(p,t)}{Q_{cW}(p,t)}, \tag{7}$$

and

$$g_{p,c}(p,t) = \frac{Q_{WM}(p,t)}{\widehat{Q}^{WM}_{0,c}(p,t)} - 1. \tag{8}$$

We report $p \in \{0.10, 0.15, \ldots, 0.90\}$. The annual summary is the unweighted mean of these 17 percentile-specific interaction estimates,[3]

$$\overline{g}_{p,c}(t) = \frac{1}{17}\sum_p g_{p,c}(p,t). \tag{9}$$

Note that this mean does not constitute a population-weighted average of individual wage premiums. Rather, it describes the interaction at selected points of the group wage distribution. The corresponding dollar summary averages wage differences over the same percentile grid.

3 Data and empirical design

3.1 IPUMS sample and wage measure

The data are from IPUMS USA, version 16.0 (Ruggles et al., 2025). The main time series uses the 1980, 1990, and 2000 decennial Censuses and ACS one-year samples from 2005 through 2024, excluding 2020 because the pandemic disrupted ACS data collection and the one-year estimates are not directly comparable with the surrounding standard releases. Appendix D adds the early ACS samples for 2000–2004 and 2020 as robustness observations. The pipeline for data pre-processing and analysis is independently implemented in Python and Julia as an additional validation step.[4]

We keep workers aged 25–64 with positive wage and salary income and positive person weights. To focus on wage rates among workers with strong labour-force attachment, the baseline sample is restricted to those who report at least 35 usual weekly hours and at least

[3]Because the pointwise benchmark need not be monotone for arbitrary wage profiles, we treat it as a rank-indexed schedule rather than assuming it is a proper quantile function.

[4]In developing and testing our analysis code, we utilized querying and auto-completion features offered by generative AI (GPT version 5.6). For further validation, we have also tasked AI agents to replicate our findings and perform stress-tests. We reviewed, revised, and validated, all code, analyses and outputs; thus, we take full responsibility for the implementation and results.

48 weeks of work. Weeks worked are available in intervals over much of the period. We divide annual wage and salary income by the midpoint of the reported weeks-worked category, using $(7, 20, 33, 43.5, 48.5, 51)$ for the six IPUMS intervals. The final category covers 50–52 weeks. Missing and not-applicable wage and salary income codes are excluded before weekly wages are constructed. Valid top-coded income observations are retained at their released values.[5]

This design produces 26,464,659 individual records across 22 main-series cross-sections. The unweighted sample ranges from 809,021 observations in 2005 to 3,910,233 in 2000. The 2024 sample contains 979,823 workers. The corresponding weighted population is 104.7 million in 2024. Table 2 in Appendix B reports sample sizes by year.

The wage measure is approximate weekly wage and salary income rather than hourly pay. This avoids imposing an exact weeks-worked measure that is unavailable in all years and keeps the wage concept consistent across the long series. For the dollar-valued results, we express wages in 2024 dollars using the CPI-U (U.S. Bureau of Labor Statistics, 2026).[6] A common deflator within a year changes dollar amounts but not $g_{p,c}$, the ratio-of-ratios, any Gini coefficient, or percentage changes in the Gini.

3.2 Race, ethnicity, and sex groups

Groups are defined from IPUMS race, Hispanic origin, and sex variables. White and Black groups are restricted to non-Hispanic respondents. Hispanic groups include Hispanic respondents of any race. The Black/White branch is

$$\mathcal{B}_B = \{WM, WW, BM, BW\}, \tag{10}$$

and the Hispanic/White branch is

$$\mathcal{B}_H = \{WM, WW, HM, HW\}. \tag{11}$$

The Hispanic category is deliberately broad to preserve a transparent four-cell comparison. It should not be read as implying homogeneity across national origin, racial identification, citizenship, or migration history. Detailed Hispanic-origin work documents substantial within-category wage differences (Ferraro et al., 2025) that our data do not allow us to decompose.

The four branch groups determine the interaction benchmark, but they do not define the population over which aggregate inequality is calculated. Every eligible worker remains in the Gini calculation, including workers whose race/ethnicity places them outside the four branch groups. In 2024, White men account for 32.3 per cent of the weighted eligible-worker sample. The Black/White branch covers 68.8 per cent of the weighted sample and the Hispanic/White branch 76.2 per cent. The remaining workers are held exactly at their observed wages in all counterfactuals.

[5]Note: Top-coding thresholds and replacement procedures vary across IPUMS samples. We consistently use the released income values in each year, without imposing additional top-code adjustments, in line with the literature working with the dataset.

[6]A common price index abstracts from differences in inflation exposure across income groups. Such differences can arise from heterogeneous consumption patterns and from the propagation of sectoral price shocks through production networks (Ipsen and Schulz, 2026). We do not attempt to construct group-specific real wage measures here.

3.3 From the rank profile to the full-sample Gini

The percentile results describe the interaction at selected points of each group distribution. Their averages describe the proportional and dollar premiums over the reported percentile grid. Separate from this, we use within-group rank-bin means (covering the entire wage distribution) to construct an aggregate inequality counterfactual, without imposing a parametric wage model. The baseline uses 20 bins. Let $\overline{Y}_{g,b,t}$ be the weighted mean wage of group g in bin b and year t. The no-interaction White-male bin mean is

$$\overline{Y}^{0,c}_{WM,b,t} = \frac{\overline{Y}_{cM,b,t}\overline{Y}_{WW,b,t}}{\overline{Y}_{cW,b,t}}. \tag{12}$$

White-male wages in that bin are multiplied by

$$s^c_{b,t} = \frac{\overline{Y}^{0,c}_{WM,b,t}}{\overline{Y}_{WM,b,t}}. \tag{13}$$

For worker i,

$$Y^{0,c}_i = \begin{cases} s^c_{b(i),t} Y_i, & i \in WM, \\ Y_i, & i \notin WM. \end{cases} \tag{14}$$

This preserves relative wage differences among White men within a rank bin while moving the bin mean to the no-interaction level. All other workers are unchanged.

Let $G(Y, a)$ denote the weighted Gini calculated from wages Y and person weights a. The Gini effect of the branch-specific no-interaction counterfactual is

$$\Delta G^0_{c,t} = G(Y, a) - G(Y^{0,c}, a). \tag{15}$$

A positive value means that the observed distribution is more unequal than the counterfactual distribution in which the White-male interaction has been removed. We report both the level difference and the percentage difference relative to the observed Gini. The applied wage adjustments implied by the counterfactual do not need to coincide with percentile-based summaries.

The counterfactual in Equation (14) lowers total wage mass whenever the White-male premium is positive.[7] To separate that feature from the distributional location of the premium, we also construct a mass-preserving benchmark. The weighted wage mass removed from White men is

$$M^c_t = \sum_{i \in WM} a_i \left(Y_i - Y^{0,c}_i\right). \tag{16}$$

We divide this amount equally, on a weighted per-person basis, among White women, minority

[7] Due to the Gini coefficient's scale invariance, a proportional redistribution of the White-male premium across all incomes would restore the original wage mass without further reducing the counterfactual Gini. Thus, any additional Gini reduction under the redistribution exercise reflects effects of allocation across workers, rather than mere mass preservation.

men, and minority women in the same branch. If $\mathcal{T}_c = \{WW, cM, cW\}$, the transfer is

$$\tau_t^c = \frac{M_t^c}{\sum_{i \in \mathcal{T}_c} a_i}. \tag{17}$$

The redistribution counterfactual is

$$Y_i^{\text{redist},c} = \begin{cases} Y_i^{0,c}, & i \in WM, \\ Y_i + \tau_t^c, & i \in \mathcal{T}_c, \\ Y_i, & \text{otherwise.} \end{cases} \tag{18}$$

This exercise is intentionally simple. It does not model labour supply, incidence, general equilibrium, or administrative feasibility. Its purpose is to show how the same intersectional wage mass affects inequality when it is not destroyed but relocated lower in the wage distribution.

3.4 Conditional wage profiles

Raw wage differences need not represent differential treatment. We therefore repeat the analysis using residualised wages. For each year, we estimate a person-weighted regression on the complete eligible-worker sample,

$$\log Y_i = \alpha_t + X_i'\beta_t + \varepsilon_i, \tag{19}$$

where X_i contains age, age squared, usual hours, education, state, metropolitan-area status, occupation, and industry. Education, state, metropolitan status, occupation, and industry enter categorically. Race-by-sex indicators are not included because the remaining group structure is precisely the object we wish to capture.

The adjusted wage index is

$$\widetilde{Y}_i = \exp\left(\overline{\log Y}_t + \widehat{\varepsilon}_i\right), \tag{20}$$

where $\overline{\log Y}_t$ is the weighted annual mean log wage. This replaces the component of log wages predicted by the included covariates with a common annual level while retaining individual residual wage differences. Exponentiation maps these residualised log wages back to a positive wage scale. The common annual level affects the scale of $\widetilde{Y}_i$ but not the interaction measures or Gini coefficients.

3.5 Inference

For ACS samples from 2005 onward, confidence intervals use the 80 person-level replicate weights distributed by IPUMS. If $\widehat{\theta}$ is the full-sample estimate and $\widehat{\theta}_r$ is the estimate from replicate r, the successive-difference-replication variance is

$$\widehat{\text{Var}}(\widehat{\theta}) = \frac{4}{80} \sum_{r=1}^{80} \left(\widehat{\theta}_r - \widehat{\theta}\right)^2. \tag{21}$$

We use a t critical value with 79 degrees of freedom for two-sided 95 per cent intervals. Replicate weights are not available for the decennial anchors or the early ACS robustness samples. For

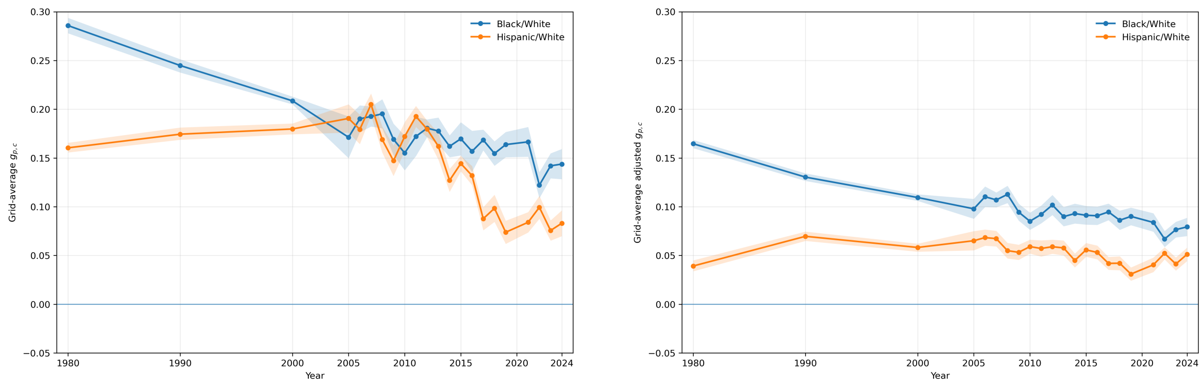

(a) Raw wages **(b)** Adjusted wage index

Fig. 2 Average excess White-male premium over time, before and after adjustment. Each panel averages $g_{p,c}(p,t)$ over the 10th–90th within-group wage percentiles. Adjusted wages are residualised within year for age, age squared, usual hours, education, state, metropolitan status, occupation, and industry. Both panels use the same vertical scale. Shaded areas are 95 per cent confidence intervals

those years, we use an exponential multiplier bootstrap at the household level. The entire adjusted-wage procedure is re-estimated within every replicate, so the adjusted intervals include uncertainty from residualisation as well as from the subsequent distributional calculations. Appendix A gives further details.

4 Results

4.1 The premium over time

The left panel of Figure 2 plots the average interaction over the 10th–90th percentile grid. The Black/White series is about 29 per cent in 1980, 25 per cent in 1990, and 21 per cent in 2000. It falls markedly over the long run but remains positive throughout the annual ACS period. The Hispanic/White series begins lower, at about 16 per cent in 1980, rises modestly through the 1990 and 2000 anchors, and then declines more strongly during the ACS period. By 2024, the estimates are 14.4 per cent for the Black/White branch and 8.3 per cent for the Hispanic/White branch.

The decline is substantively important. The historical evidence does not suggest that the interaction is a fixed proportion of wages. The Black/White premium has roughly halved since 1980, and the Hispanic/White premium is now well below its 1990s level. Yet neither series approaches zero. In that sense, the long-run development combines convergence with persistence: the nonadditive component has narrowed while remaining economically meaningful. This is a different statement from the post-2000 widening of gender-specific racial pay gaps documented by Mandel and Semyonov (2016). Our interaction can narrow even when the underlying racial gaps remain large, or widen, if the gaps among men and women become more similar. What falls in the left panel of Figure 2 is therefore the nonadditivity between the two gaps, not the racial pay gap itself.

The left panel of Figure 3 shows where the 2024 interaction sits in the wage distribution. For

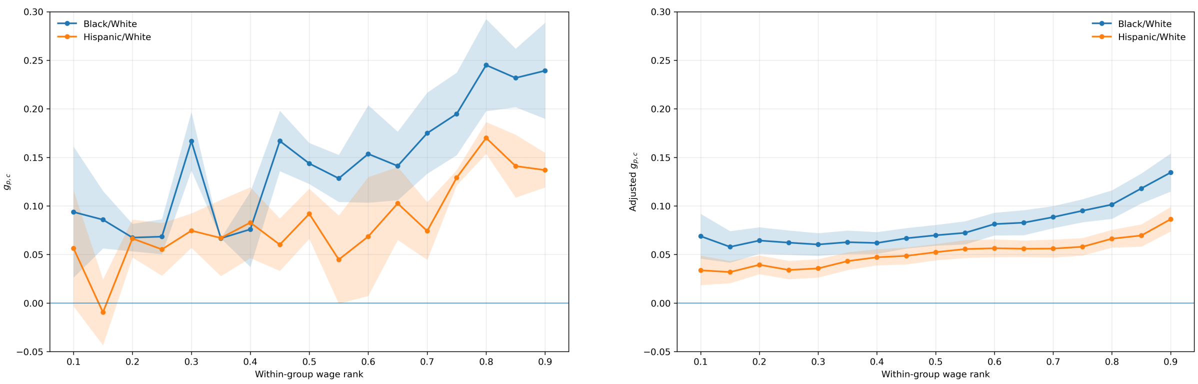

(a) Raw wages

(b) Adjusted wage index

Fig. 3 Excess White-male premium by within-group wage rank, 2024, before and after adjustment. The premium is measured relative to the no-interaction benchmark in Equation (4). Adjusted wages are residualised within year for age, age squared, usual hours, education, state, metropolitan status, occupation, and industry. Both panels use the same vertical scale. Shaded areas are 95 per cent confidence intervals

the Black/White comparison the point estimate is positive at every reported percentile. It is around 9 per cent at the 10th percentile, fluctuates between roughly 7 and 17 per cent through the middle of the distribution, and then rises steeply: about 20 per cent at the 75th percentile and about 24 per cent at the 90th. The Hispanic/White profile is smaller and noisier at the bottom. Its 15th-percentile estimate is slightly negative, and some lower-middle confidence intervals include zero. From the upper-middle part of the distribution onward, the premium is clearly positive and reaches roughly 14 per cent at the 90th percentile.

The upper-tail gradient is important for two reasons. It shows that the average series in the left panel of Figure 2 does not arise from a constant proportional shift in White-male wages, and it places much of the associated wage mass among workers who already occupy high positions in the pooled distribution. Appendix E reports the equivalent ratio-of-ratios $1 + g_{p,c}$.

4.2 Dollar magnitudes and aggregate inequality

Figure 4 converts the proportional interaction into real weekly wage amounts. The dollar profile is much steeper than the proportional profile because wage levels themselves increase with rank. In the Black/White branch, the estimated premium rises from roughly \$60 per week near the bottom of the reported range to more than \$700 at the 90th percentile. The Hispanic/White amount rises from near zero in the lower tail to about \$450 at the 90th percentile. Averaged over the full percentile grid, the 2024 amounts are \$243 and \$152 per week.

The consequence for the full wage distribution is shown in Figure 5. The plotted quantity is the percentage reduction in the Gini when White-male wages are moved to the branch-specific no-interaction benchmark. The Gini effect is positive in every year in both branches. It is largest in the decennial anchors, especially for the Black/White branch, and declines over time along with the premium. In 2024, the Black/White counterfactual lowers the Gini by 3.17 per cent (95 per cent CI 2.50–3.85), while the Hispanic/White counterfactual lowers it by 2.80 per cent (95 per cent CI 2.27–3.34).

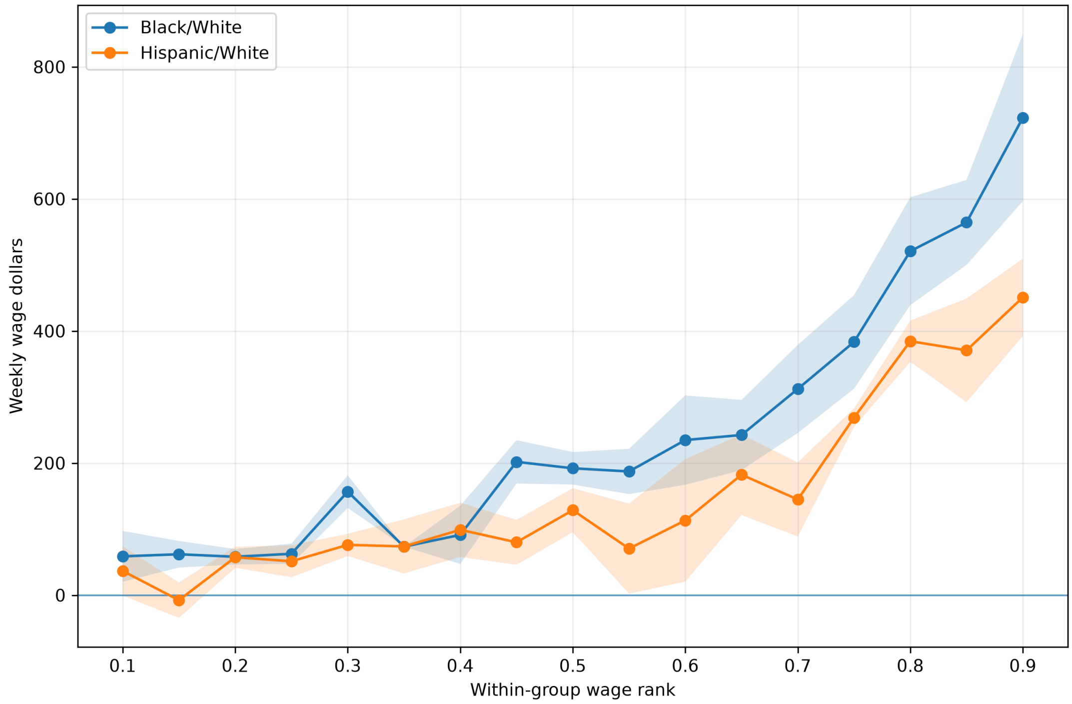

Fig. 4 White-male excess weekly wage premium by within-group wage rank, 2024, in real 2024 dollars. Shaded areas are 95 per cent confidence intervals

These changes are large enough to matter in the context of aggregate wage inequality. In 2024, the actual weighted Gini is 0.3997. The Black/White benchmark lowers it to 0.3871 and the Hispanic/White benchmark to 0.3885. The mass-preserving exercise sharpens this interpretation. Figure 6 compares the benchmark counterfactual with the redistribution counterfactual. In every year, the redistribution effect is larger. In 2024, returning the Black/White interaction wage mass as an equal per-person transfer to White women, Black men, and Black women lowers the Gini by 8.33 per cent (95 per cent CI 7.38–9.28). The corresponding Hispanic/White exercise lowers it by 7.44 per cent (95 per cent CI 6.51–8.37). The thought experiment thus illustrates that the location of the wage mass, not merely its existence, is central to its inequality effect.

Table 1 collects the latest-year estimates. The adjusted results discussed below make clear that the raw interaction is partly associated with observed worker and job characteristics, but they do not eliminate either the premium or its aggregate inequality consequence.

4.3 Conditional wage structure

Adjustment for standard wage-gap covariates substantially reduces the premium but does not remove it. The right panel of Figure 2 shows the annual average for the residualised wage index. The Black/White estimate declines from about 16.5 per cent in 1980 to 7.9 per cent in 2024. The Hispanic/White estimate is smaller throughout most of the period and is 5.1 per cent in 2024. Both estimates are significantly positive in the latest year.

The right panel of Figure 3 shows that the adjusted rank profile is smoother than the raw profile. Both branches are positive across the entire reported range. The Black/White premium is approximately 6–7 per cent through much of the lower half and rises to 13.5 per cent at

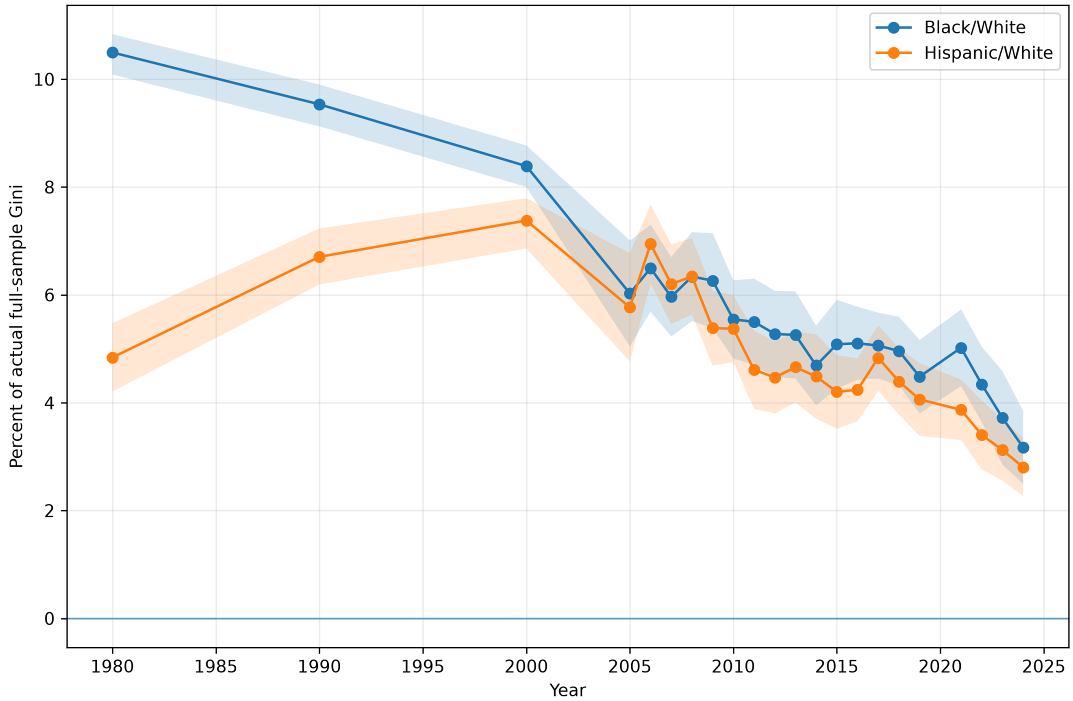

Fig. 5 Percentage reduction in the Gini of the complete eligible-worker sample when White men are moved to the no-interaction benchmark. Workers outside the four branch groups remain in the distribution and are left unchanged. Shaded areas are 95 per cent confidence intervals

the 90th percentile. The Hispanic/White premium increases from roughly 3 per cent near the bottom to 8.6 per cent at the 90th percentile. The fact that the upper-tail gradient survives the adjustment is particularly relevant for inequality: it is not solely generated by the measured occupational, industrial, educational, or geographic composition of White-male employment.

The adjusted Gini results are shown in Figure 7. The adjusted wage index has a 2024 Gini of 0.2898. Moving White men to the Black/White benchmark reduces it to 0.2801, a 3.32 per cent decline; the Hispanic/White benchmark reduces it to 0.2828, a 2.40 per cent decline. The level changes in the Gini are smaller than for raw wages, but the percentage changes remain substantial because the residualised wage distribution is less dispersed to begin with.

The conditional analysis should not be read as a causal specification. Some controls may themselves mediate structural advantage, and a common-slope regression imposes a particular adjustment technology. Its role here is narrower, namely, to show that the raw results are not driven only by easily observed differences in age, education, hours, location, occupation, or industry, and that the remaining interaction still has a measurable relationship with aggregate wage inequality.

5 Interpretation and scope

The four observed wage profiles identify the interaction in Equation (6). They do not identify which social process generated it, nor do they uniquely determine which of the four cells should be altered in a counterfactual. Our convention holds White women, minority men, and minority women fixed and assigns the nonadditive component to White men. This is the appropriate

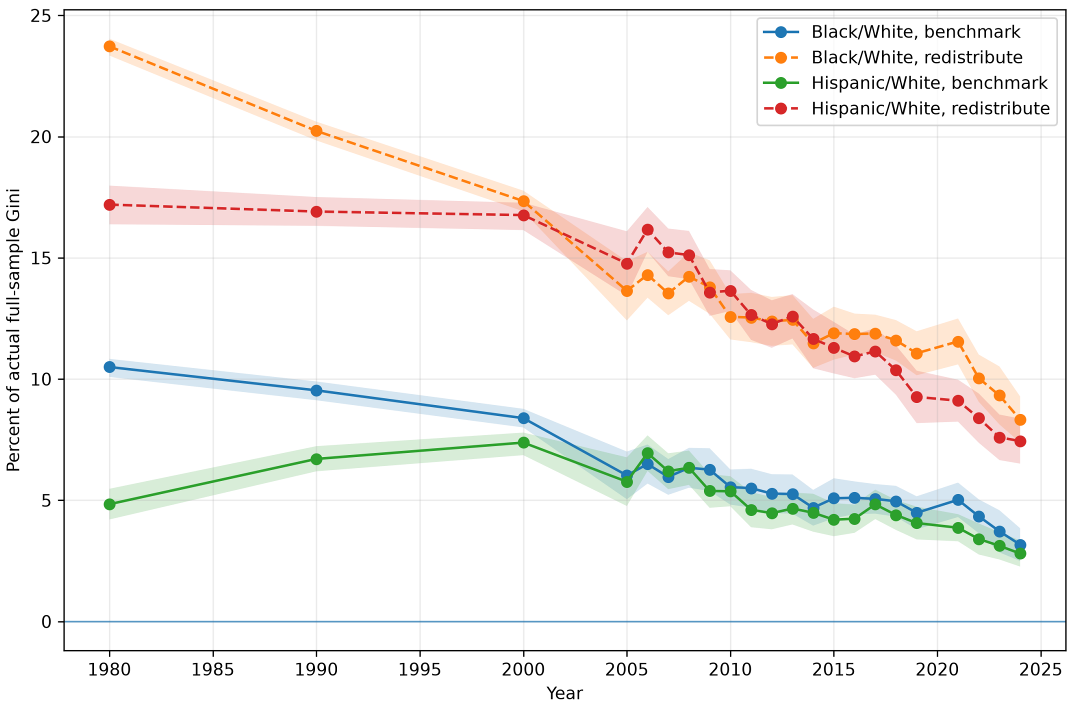

Fig. 6 Percentage reduction in the full-sample Gini under the no-interaction benchmark and the mass-preserving redistribution counterfactual. Shaded areas are 95 per cent confidence intervals

convention for the question posed here–how much inequality is associated with the interaction if it is represented as excess privilege–but it is not a causal decomposition of individual wages.

This distinction helps reconcile the paper with the broader intersectionality literature. Studies that estimate an additional penalty for Black women are asking whether multiply marginalised workers fall below an additive benchmark (George et al., 2022; Paul et al., 2022). Other studies find little amplifying interaction or a smaller-than-additive disadvantage (Greenman and Xie, 2008; McLaughlin and Neumark, 2026). Our results do not imply that one set of findings is generally correct and the other is generally wrong. They show that, in the IPUMS wage distributions examined here, a privilege-oriented multiplicative benchmark yields a positive average interaction, especially at high wage ranks, and that this component has a nontrivial aggregate inequality effect.

Selection is an important limitation. We condition on being a full-time, full-year wage earner for data availability reasons. That is a substantively meaningful population for studying wage rates, as operationalized by weekly earnings, but it is not the entire working-age population. If race, gender, employment, and potential wages jointly affect selection into the analytic sample, observed interactions need not equal population interactions. McLaughlin and Neumark (2026) explicitly examine employment selection in intersectional wage comparisons, and Bashir and Al-Kassab-Córdova (2026) show more generally that conditioning on selected samples can create, attenuate, or reverse interaction patterns. Our estimates should therefore be interpreted as describing wage inequality among workers with strong labour-force attachment. They do not quantify intersectional inequality in employment probabilities or in total annual resources (including also capital income): These factors might be partially captured by the wage differences

Table 1: Headline results, 2024

Outcome	Branch	Avg. premium (percent)	Weekly premium (2024 dollars)	Gini reduction (percent)	Redistribution (percent)
Raw	Black/White	14.38	242.74	3.17	8.33
Raw	Hispanic/White	8.31	151.90	2.80	7.44
Adjusted	Black/White	7.93	119.25	3.32	6.55
Adjusted	Hispanic/White	5.12	79.67	2.40	4.83

Notes: The average proportional and weekly dollar premium refer to the mean of $g_{p,c}(p, 2024)$ over the 10th–90th percentile grid; they do not measure the average individual adjustment in the full-sample counterfactual. The Gini reduction moves White men to the no-interaction benchmark and leaves all other wages unchanged. Redistribution transfers the removed weighted wage mass equally per recipient among the other three branch groups. All Ginis are calculated on the complete eligible-worker sample.

that we analyze but also constitute separate manifestations of discrimination and (excessive While male) privilege that we invite future research to study though it might be harder to capture quantitatively than the wage data presented here.

The two branches are also intentionally coarse. The Black/White comparison omits heterogeneity within the categories it uses, and the Hispanic/White comparison combines people with very different national origins and migration histories. The detailed Hispanic evidence in Ferraro et al. (2025) demonstrates the relevance of this issue. The advantage of the four-cell design is transparency: the interaction and its counterfactual are easy to interpret. The cost is that it does not represent the full multidimensionality of U.S. stratification. Extending the inequality accounting to more detailed groups is feasible, but it requires additional normative choices about which interactions are assigned to which cells. More importantly, the analysed groups naturally get much smaller as the number of groups increases, limiting statistical power.

Policy implications are necessarily more tentative because the estimated premium is descriptive rather than causal. De Poli and Maier (2025) study the distributional consequences of narrowing gender pay penalties in the context of the European Union's pay-transparency agenda. The same transparency logic could directly be extended beyond one-dimensional gender gaps. Where the relevant data are available, reporting requirements could include intersectional contrasts, such as differences in the gender pay gap across racial or ethnic groups, and thereby make a nonadditive premium of the kind measured here visible alongside the conventional gender pay gap. The estimates presented here suggest that transparency restricted to one dimension can leave an aggregate distributional component unreported.

Finally, the redistribution counterfactual is an accounting benchmark. It preserves total wage mass and demonstrates the inequality consequences of moving the interaction-associated wage amount from a high-wage group to lower-wage groups. A real intervention could change labour supply, occupational choice, bargaining, firm behaviour, taxes, prices, and the wage structure itself. Those behavioural and general-equilibrium responsesare outside the present design. The observed nonadditive wage structure is sufficiently concentrated among White men, and sufficiently high in their wage distribution, that removing it changes the overall Gini by several per cent.

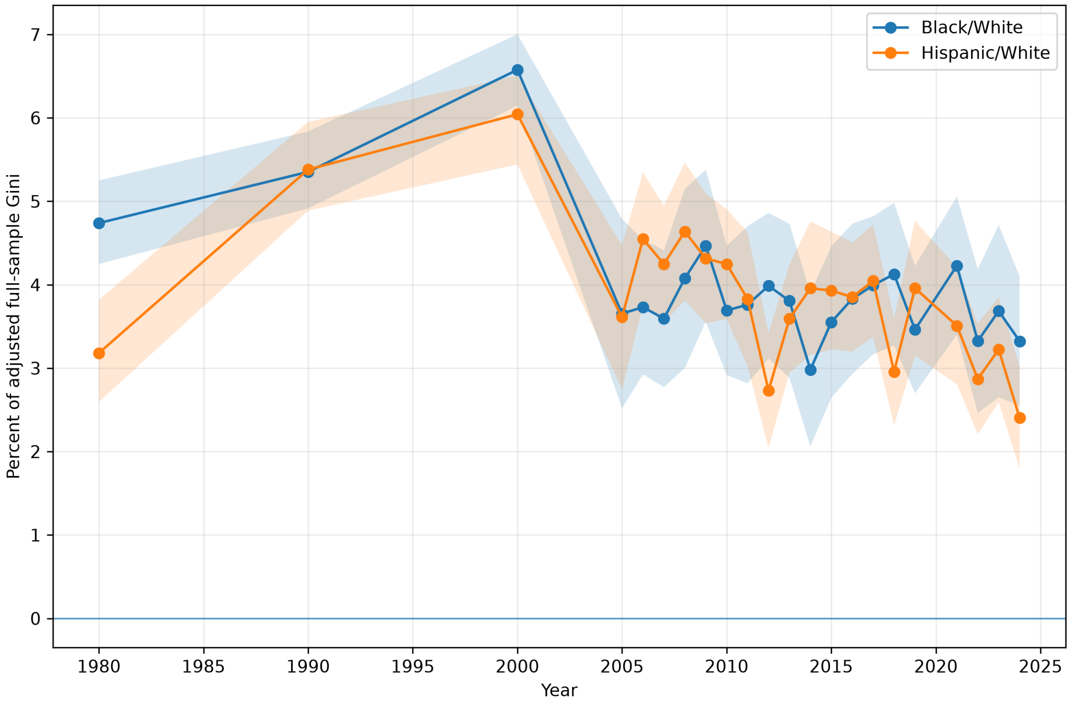

Fig. 7 Percentage reduction in the Gini of the adjusted complete eligible-worker sample when White men are moved to the no-interaction benchmark. Shaded areas are 95 per cent confidence intervals

6 Conclusion

Intersectional wage analysis is usually organised around disadvantage. That focus is natural, but it can obscure the distributional role of advantage. A four-group wage structure contains a single nonadditive component, not separate interaction terms for disadvantage and privilege. Its interpretation depends on the allocation convention: the same departure from additivity can be represented by adjusting the multiply marginalised cell or, equivalently, by adjusting the multiply advantaged cell. Our analysis adopts the latter representation and asks what follows for aggregate inequality when the interaction is allocated to White men.

For U.S. full-time, full-year workers, the privilege-oriented interaction is positive on average in both the Black/White and Hispanic/White comparisons from 1980 through 2024. It has declined over the long run, but remains large in recent data. In 2024, the average premium is 14.4 per cent for the Black/White branch and 8.3 per cent for the Hispanic/White branch, and it rises strongly toward the upper part of the wage distribution. Conditional estimates are smaller but preserve the same basic pattern.

The paper's main contribution is the link from this intersectional structure to inequality in the entire wage distribution. Moving White men to the no-interaction benchmark lowers the 2024 full-sample Gini by 3.17 per cent in the Black/White comparison and 2.80 per cent in the Hispanic/White comparison. Keeping total wage mass fixed and reallocating the removed amount to the other branch groups produces still larger Gini reductions. These effects arise because the interaction is attached to a large population group and because much of its dollar value is located at high pooled wage ranks.

This shifts the interpretation of intersectionality toward the distributional consequences of nonadditive wage structure. Race-by-sex nonadditivity is one component of aggregate wage inequality, and its importance depends both on the size of the interaction and on where the associated wage mass is located in the distribution. A natural next step is therefore to identify the mechanisms that generate this pattern. Future work could examine, for example, the roles of occupational sorting, firm assignment, bargaining, work organisation, and differential returns to worker and job characteristics, and ask which of these mechanisms account for the particularly pronounced premium at higher wage ranks. Distinguishing among these channels would help move from the descriptive inequality accounting developed here toward a causal understanding of how intersectional advantage is produced and sustained.

A Survey-design and bootstrap inference

For ACS samples from 2005 onward, all raw and adjusted statistics are recomputed with each of the 80 person-level replicate weights. If $\widehat{\theta}_r$ denotes the estimate from replicate r, we use

$$\widehat{\mathrm{Var}}(\widehat{\theta}) = \frac{4}{80} \sum_{r=1}^{80} \left(\widehat{\theta}_r - \widehat{\theta}\right)^2 . \tag{22}$$

The 95 per cent intervals use a t_{79} critical value. For the 1980, 1990, and 2000 Census anchors and the 2000–2004 ACS robustness samples, we use an exponential multiplier bootstrap clustered at the household level. For bootstrap draw m,

$$a_i^{(m)} = a_i \eta_{h(i)}^{(m)}, \qquad \eta_h^{(m)} \sim \mathrm{Exp}(1), \tag{23}$$

where $h(i)$ indexes the household. The reported intervals for these samples are percentile intervals.

For adjusted outcomes, the wage regression is re-estimated in every replicate. Adjusted wages, weighted quantiles, rank-bin means, White-male scaling factors, and Gini coefficients are then reconstructed from that replicate. The uncertainty bands therefore reflect the complete estimator rather than treating the first-stage residualisation as fixed.

B Sample sizes

Table 2: Analytic sample size by main-series cross-section

Year	Observations	Weighted population (millions)
1980	2,529,862	50.60
1990	3,248,554	66.03
2000	3,910,233	79.95
2005	809,021	82.93
2006	827,491	84.86
2007	834,872	85.93
2008	888,213	90.98
2009	851,244	86.84
2010	835,854	85.03
2011	819,863	86.02
2012	834,795	87.74
2013	855,922	88.85
2014	861,365	90.71
2015	878,396	92.84
2016	889,860	94.19
2017	913,394	96.45
2018	927,253	97.94
2019	940,436	99.42
2021	887,098	95.97
2022	963,402	101.65
2023	977,708	103.17
2024	979,823	104.68

C Rank-bin sensitivity

The aggregate counterfactual uses 20 within-group rank bins in the baseline. With one bin, the White-male adjustment is a single mean-level scaling. Increasing the number of bins allows the counterfactual to follow more of the observed distributional shape. Figures 8 and 9 report the 2024 Gini results for alternative bin counts. The headline conclusion is not driven by the baseline choice.

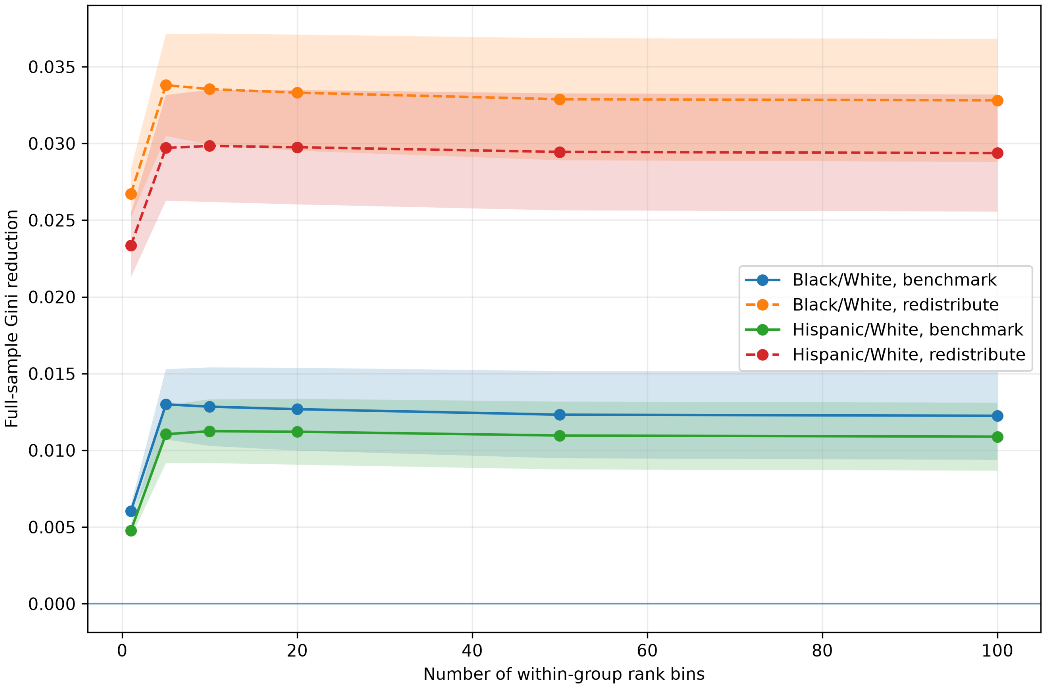

Fig. 8 Sensitivity of the raw 2024 full-sample Gini counterfactual to the number of within-group rank bins. Shaded areas are 95 per cent confidence intervals

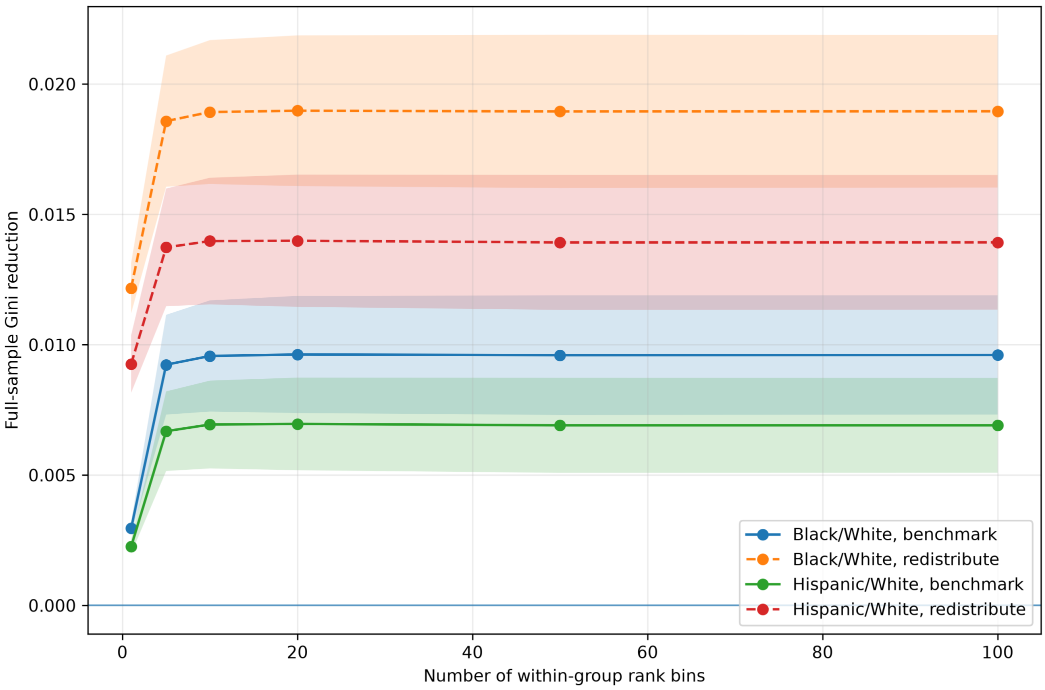

Fig. 9 Sensitivity of the adjusted 2024 full-sample Gini counterfactual to the number of within-group rank bins. Shaded areas are 95 per cent confidence intervals

D Early ACS years and 2020

The main series uses decennial Census anchors and ACS samples from 2005 onward and omits 2020. Figures 10–13 add the early ACS one-year samples from 2000–2004 and the 2020 observation. Because the 2000 Census and 2000 ACS refer to the same calendar year, the robustness series uses the 2000 ACS value in place of the Census anchor. Hollow markers denote supplemental observations.

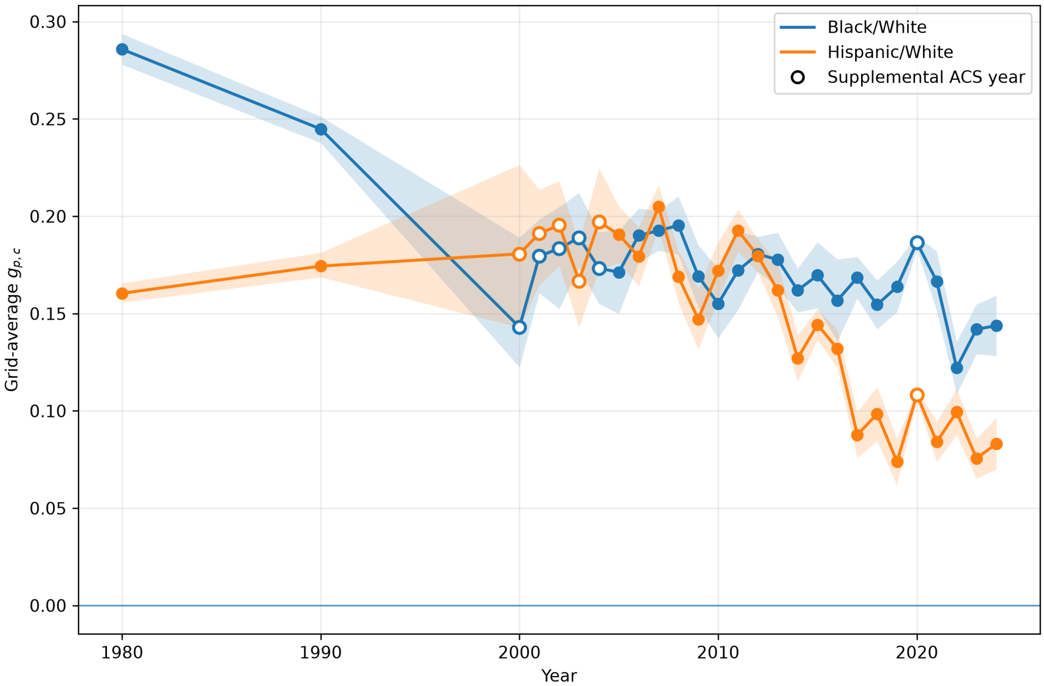

Fig. 10 Average excess White-male premium in the extended ACS series. Hollow markers indicate the early ACS years and 2020

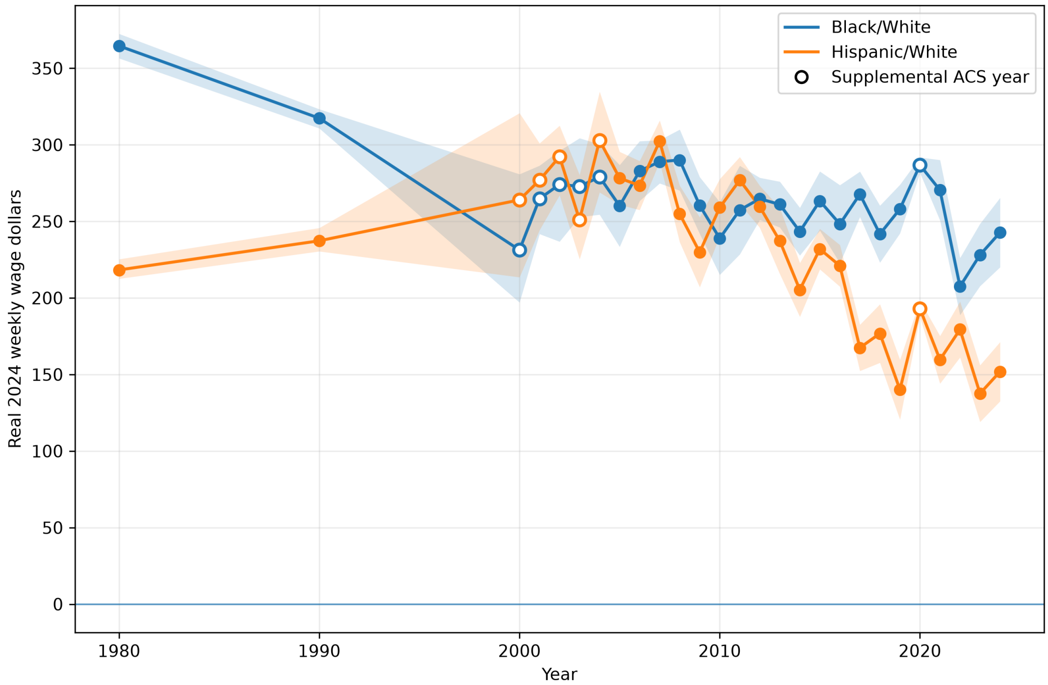

Fig. 11 Average excess White-male weekly wage premium in real 2024 dollars in the extended ACS series. Hollow markers indicate the early ACS years and 2020

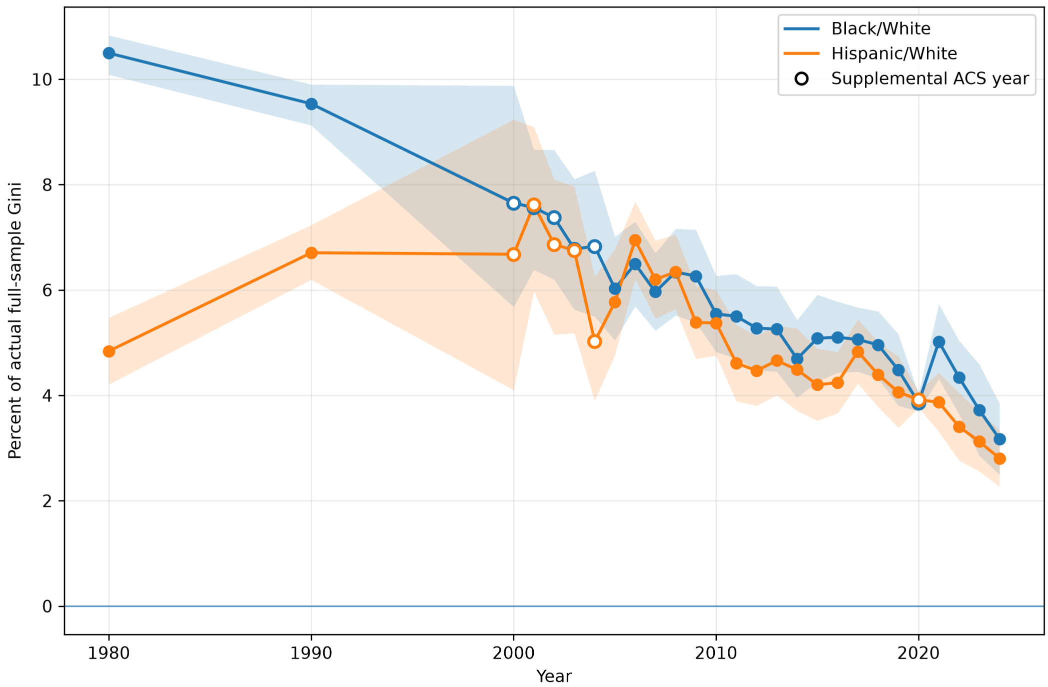

Fig. 12 Percentage reduction in the full-sample Gini under the no-interaction counterfactual in the extended ACS series. Hollow markers indicate the early ACS years and 2020

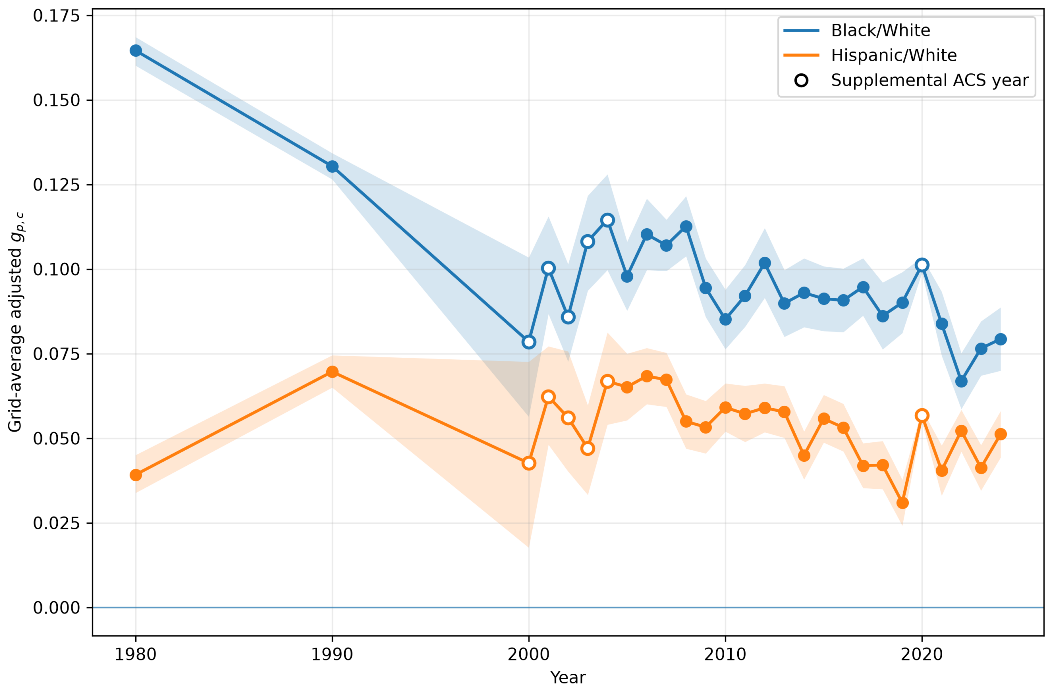

Fig. 13 Adjusted average excess White-male premium in the extended ACS series. Hollow markers indicate the early ACS years and 2020

E Additional figures

Figures 14–17 report the ratio-of-ratios, the average real-dollar premium, the level change in the Gini, and the adjusted redistribution exercise.

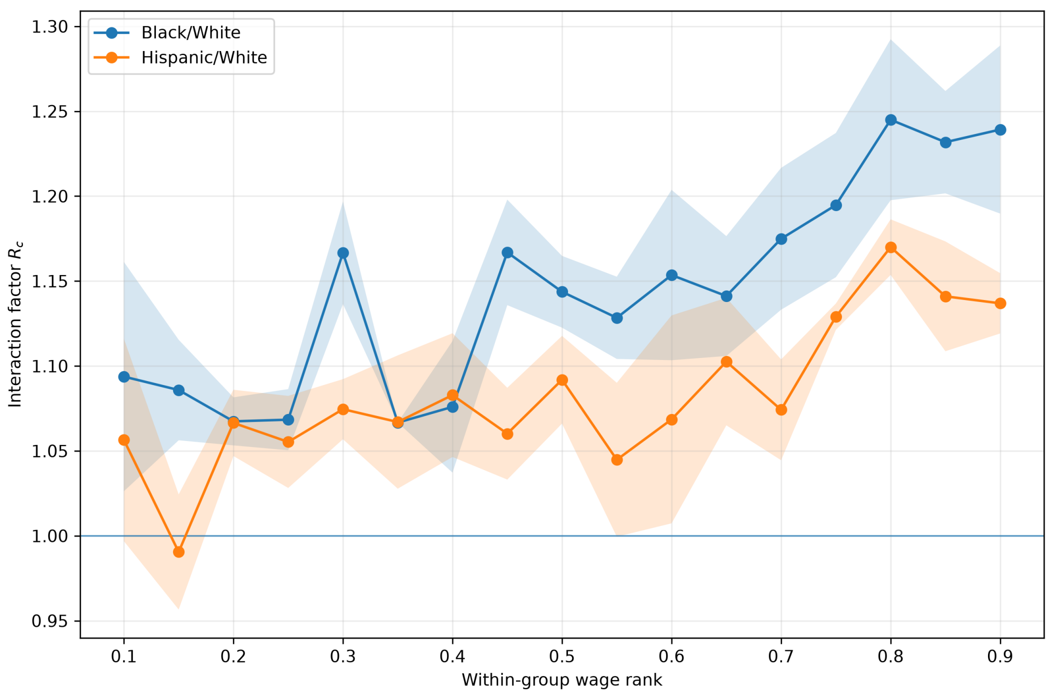

Fig. 14 Ratio-of-ratios $1 + g_{p,c}$ by within-group wage rank, 2024

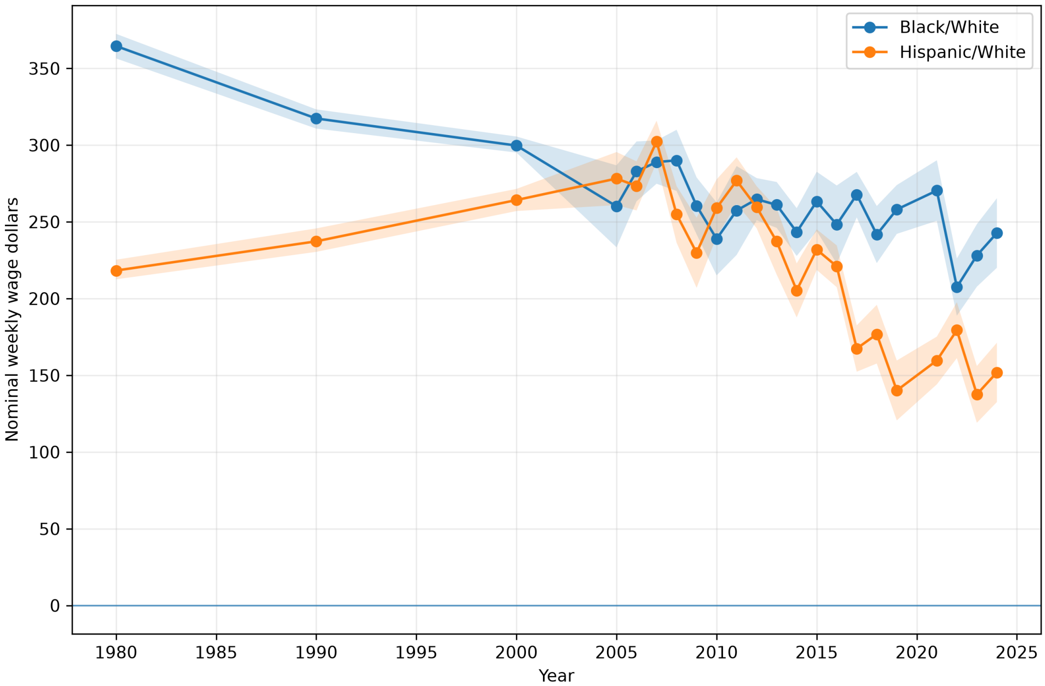

Fig. 15 Average White-male excess weekly wage premium over time in real 2024 dollars

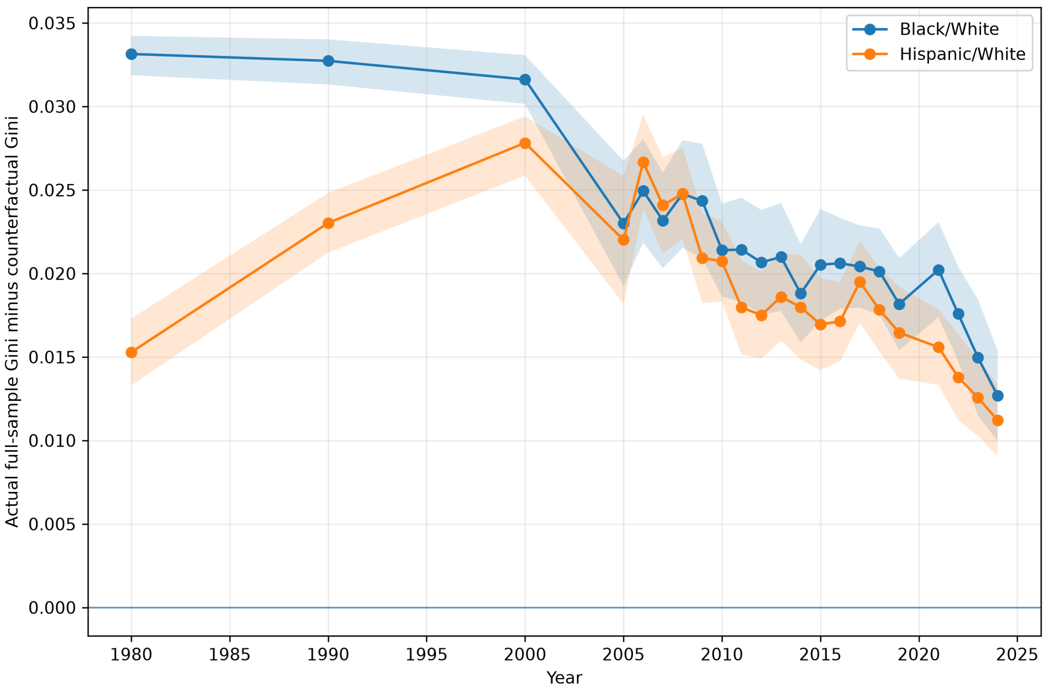

Fig. 16 Level reduction in the full-sample Gini when White men are moved to the no-interaction benchmark

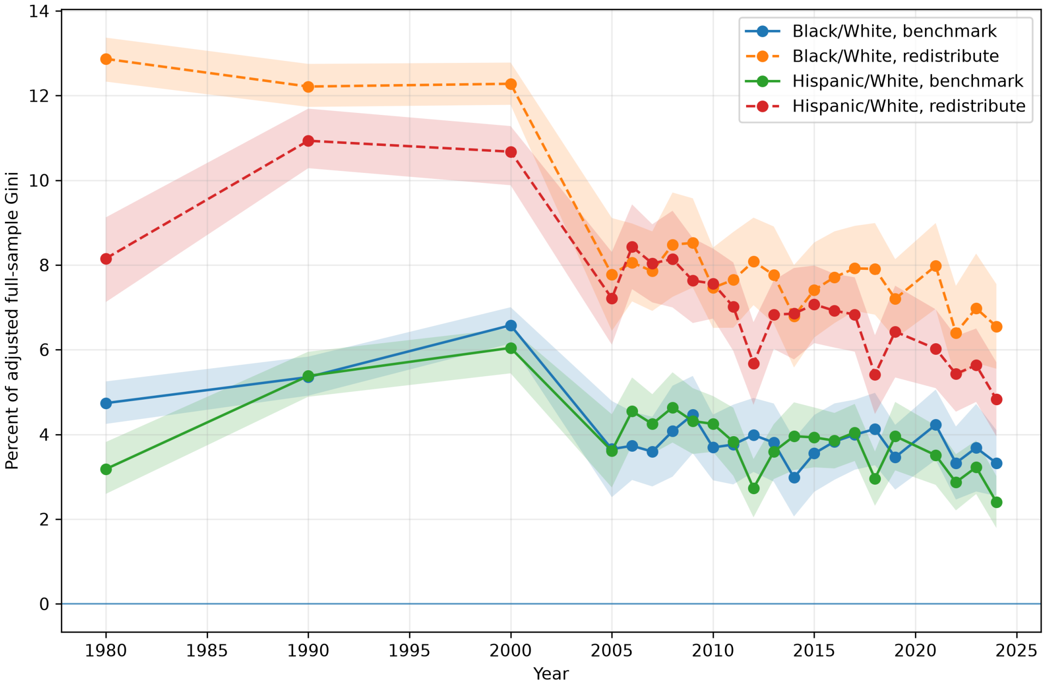

Fig. 17 Adjusted full-sample Gini reductions under the no-interaction and mass-preserving redistribution counterfactuals

F When does the premium raise the Gini?

The counterfactual results above are computed directly. To characterise when the estimated premium increases overall wage inequality, introduce a scaling parameter $t \in [0, 1]$ that determines how much of the premium is added to the no-interaction wage distribution. Let Y_i^0 be worker i's no-interaction wage and define $D_i = Y_i - Y_i^0$ for White men and $D_i = 0$ for all other workers. Define

$$Y_i(t) = Y_i^0 + tD_i, \qquad t \in [0, 1]. \tag{24}$$

At $t = 0$, Equation (24) gives the no-interaction counterfactual, while at $t = 1$ it reproduces the observed wage distribution. Intermediate values of t provide a mathematical device for introducing the estimated premium gradually and examining its marginal effect on inequality.

Let w_i be normalised person weights, $\mu(t)$ the pooled mean wage, $G(t)$ the pooled Gini, and F_t^m the pooled mid-distribution function associated with the wage distribution at t. For values of t at which the derivative exists,

$$\frac{\partial G(t)}{\partial t} = \frac{2}{\mu(t)} \sum_{i \in WM} w_i D_i \left[F_t^m \left(Y_i(t) \right) - \frac{1 + G(t)}{2} \right]. \tag{25}$$

Equation (25) follows by differentiating the mean absolute difference representation of the weighted Gini. The derivative has a local interpretation: at a given value of t, it measures whether adding a small additional share of the estimated premium raises or lowers the Gini.

The term $F_t^m(Y_i(t))$ records worker i's position in the pooled wage distribution. The mid-distribution function assigns observations tied at the same wage their midpoint rank, which is the appropriate rank convention for the derivative of the weighted Gini. The subscript t is important because adding the premium can change workers' relative positions in the pooled distribution. A White-male worker may therefore occupy a different pooled rank at $t = 0$ than at $t = 1$.

Suppose now that $D_i \geq 0$ for all White-male workers and that the total premium is positive. Define the premium-weighted pooled rank as

$$\overline{R}^{D}(t) = \frac{\sum_{i \in WM} w_i D_i F_t^m \left(Y_i(t)\right)}{\sum_{i \in WM} w_i D_i}. \tag{26}$$

Because the weights $w_i D_i$ are non-negative, $\overline{R}^{D}(t)$ has a direct interpretation: it is the average pooled rank at which the wage mass generated by the premium is received. Thus, it differs from the average rank of White men. Workers receiving larger premiums receive correspondingly greater weight.

Using Equation (26), Equation (25) can be written as

$$\frac{\partial G(t)}{\partial t} = \frac{2}{\mu(t)} \left(\sum_{i \in WM} w_i D_i \right) \left[\overline{R}^{D}(t) - \frac{1 + G(t)}{2} \right]. \tag{27}$$

It follows immediately that

$$\frac{\partial G(t)}{\partial t} > 0 \quad \Longleftrightarrow \quad \overline{R}^{D}(t) > \frac{1 + G(t)}{2}. \tag{28}$$

Equation (28) shows that a positive premium does not mechanically increase overall wage inequality. Its effect also depends on where the associated wage mass lies in the pooled wage distribution. The premium raises the Gini when it is sufficiently concentrated among workers who already occupy high pooled ranks. Conversely, a positive premium concentrated sufficiently far down the distribution can reduce the Gini by compressing wage differences.

The relevant threshold is not the median. Because the Gini is normalised by mean income, adding wage mass also raises $\mu(t)$. The additional wage mass must therefore be located sufficiently high in the distribution for its effect on wage dispersion to dominate this normalisation effect. For example, if $G(t) = 0.4$, Equation (28) places the threshold at 0.7. At that value of t, the premium raises the Gini locally only if its premium-weighted average pooled rank exceeds the 70th percentile.

Finally, Equation (28) is a local condition. The exact difference between the observed and no-interaction Ginis is obtained by accumulating these marginal effects over the full scaling interval:

$$G(1) - G(0) = \int_0^1 \frac{\partial G(t)}{\partial t} \, dt. \tag{29}$$

If the derivative in Equation (25) is positive throughout the interval, apart from values at which the derivative is not defined, the observed Gini must exceed the no-interaction Gini. If its sign changes as the premium is scaled up, the overall effect is instead determined by the integral in Equation (29). Non-differentiability can occur when workers exchange ranks or when wage ties are created or broken; these points do not affect the integral comparison between $G(0)$ and $G(1)$.

Statements and Declarations

Author contributions **Conceptualization:** Jan Schulz, Caleb E. Agoha, Daniel M. Mayerhoffer; **Methodology:** Jan Schulz; **Formal analysis:** Jan Schulz; **Investigation:** Caleb E. Agoha (literature review); Jan Schulz **Software:** Jan Schulz, Daniel M. Mayerhoffer; **Writing – original draft:** Jan Schulz, Caleb E. Agoha, Daniel M. Mayerhoffer; **Writing – review and editing:** Jan Schulz, Caleb E. Agoha, Daniel M. Mayerhoffer; **Project administration:** Jan Schulz.

Funding The authors did not receive external funding from any organisation for the submitted work. Mayerhoffer's contribution is supported by the University of Amsterdam Data Science Centre.

Competing interests The authors have no relevant financial or non-financial interests to disclose.

Data availability The IPUMS USA microdata used in this study are available to registered users from IPUMS USA (Ruggles et al., 2025), subject to the applicable terms of use. The raw IPUMS microdata cannot be redistributed by the authors. CPI-U data are publicly available from the U.S. Bureau of Labor Statistics. Access information and a full list of data sources will be included with the replication materials.

Code availability Replication code and documentation required to reproduce the analysis are publicly available at `https://anonymous.4open.science/r/excessive-white-male-privilege-in-data-1` (Will be replaced by the direct link to the github repository for publication.)

Acknowledgements We are grateful to Daniel Hilbinger for helpful research assistance and thank participants at the 2026 Conference of the European PPE Network for helpful comments and discussion. All remaining errors are our own.

References

Almquist, E. M. (1975). Untangling the Effects of Race and Sex: The Disadvantaged Status of Black Women. *Social Science Quarterly*, 56(1):129–142.

Alonso-Villar, O. and del Río, C. (2023). Privilege and hindrance on the USA earnings distribution by gender and race/ethnicity: An intersectional framework with 12 groups. *International Journal of Manpower*, 44(4):635–652. `https://doi.org/10.1108/IJM-12-2021-0705`.

Alonso-Villar, O. and del Río, C. (2024). Gender, race, and class in an intersectional framework: Occupations and wages in the united states. *Feminist Economics*, 30(1):40–69. `https://doi.org/10.1080/13545701.2023.2255871`.

Bashir, N. Z. and Al-Kassab-Córdova, A. (2026). Is intersectionality selective? the role of collider bias. *Social Science & Medicine*, 404:119550. `https://doi.org/10.1016/j.socscimed.2026.119550`.

Budig, M. J., Lim, M., and Hodges, M. J. (2021). Racial and gender pay disparities: The role of education. *Social Science Research*, 98:102580. https://doi.org/10.1016/j.ssresearch.2021.102580.

Cech, E. A. (2022). The intersectional privilege of white able-bodied heterosexual men in STEM. *Science Advances*, 8(24):eabo1558. https://doi.org/10.1126/sciadv.abo1558.

Cowell, F. A. and Jenkins, S. P. (1995). How much inequality can we explain? a methodology and an application to the united states. *The Economic Journal*, 105(429):421–430. https://doi.org/10.2307/2235501.

Crenshaw, K. (1989). Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *University of Chicago Legal Forum*, 1989(1):139–167.

De Poli, S. and Maier, S. (2025). Enforcing 'equal pay for equal work' in the EU: What would it take? *The Journal of Economic Inequality*, 23:1219–1254. https://doi.org/10.1007/s10888-024-09654-5.

del Río, C., Gradín, C., and Cantó, O. (2011). The measurement of gender wage discrimination: The distributional approach revisited. *The Journal of Economic Inequality*, 9:57–86. https://doi.org/10.1007/s10888-010-9130-7.

Ferraro, A. C., Young, A., Bernal, B., Vo, T., and Schleifer, C. (2025). Hispanic intersectional pay in the united states: Measuring the intra- and inter-ethnoracial and gender income differences by hispanic groups. *Sociological Forum*, 40(3):319–341. https://doi.org/10.1111/socf.13046.

George, E. E., Milli, J., and Tripp, S. (2022). Worse than a double whammy: The intersectional causes of wage inequality between women of colour and white men over time. *LABOUR*, 36(3):302–341. https://doi.org/10.1111/labr.12226.

Giangregorio, L. and Villani, D. (2025). Functional Distribution, Personal Income Inequality, and Top Shares of Income: Do Social Classes Still Matter? *The Journal of Economic Inequality*, 23:143–176. https://doi.org/10.1007/s10888-024-09633-w.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119. https://doi.org/10.1257/aer.104.4.1091.

Greenman, E. and Xie, Y. (2008). Double jeopardy? the interaction of gender and race on earnings in the united states. *Social Forces*, 86(3):1217–1244. https://doi.org/10.1353/sof.0.0008.

Heikkuri, V.-M. and Schief, M. (2026). Subgroup decomposition of the gini coefficient: A new solution to an old problem. *Econometrica*, 94(1):169–192. https://doi.org/10.3982/ECTA22145.

Ipsen, L. and Schulz, J. (2026). The (dis-)equalizing effects of production networks. *Applied Economics Letters*, 33(6):830–834. https://doi.org/10.1080/13504851.2024.2396557.

Kim, M. (2009). Race and gender differences in the earnings of black workers. *Industrial Relations: A Journal of Economy and Society*, 48(3):466–488. https://doi.org/10.1111/j.1468-232X.2009.00569.x.

Madsen, S., Medina, R., Brewer, S., and Linke, A. (2026). Intersectionality in the united states: Examining impacts of identity using US census microdata. *GeoJournal*, 91:123. https://doi.org/10.1007/s10708-026-11681-w.

Mandel, H. and Semyonov, M. (2016). Going back in time? gender differences in trends and sources of the racial pay gap, 1970 to 2010. *American Sociological Review*, 81(5):1039–1068. https://doi.org/10.1177/0003122416662958.

McKinney, K. L., Abowd, J. M., and Janicki, H. P. (2022). U.S. long-term earnings outcomes by sex, race, ethnicity, and place of birth. *Quantitative Economics*, 13(4):1879–1945. https://doi.org/10.3982/QE1908.

McLaughlin, J. S. and Neumark, D. (2026). Is there intersectional labor market discrimination? *Industrial Relations: A Journal of Economy and Society*. Advance online publication; https://doi.org/10.1111/irel.70033.

Paul, M., Zaw, K., and Darity, W. (2022). Returns in the labor market: A nuanced view of penalties at the intersection of race and gender in the US. *Feminist Economics*, 28(2):1–31. https://doi.org/10.1080/13545701.2022.2042472.

Reiss, L. K., Schiffinger, M., Rapp, M. L., and Mayrhofer, W. (2025). Intersectional income inequality: A longitudinal study of class and gender effects on careers. *Culture and Organization*, 31(6):473–492. https://doi.org/10.1080/14759551.2023.2232505.

Restifo, S. J., Ryabov, I., and Ruiz, B. (2023). Race, gender, and nativity in the southwest economy: An intersectional approach to income inequality. *Population Research and Policy Review*, 42:48. https://doi.org/10.1007/s11113-023-09779-x.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J. A., Richards, S., Rodgers, R., Schroeder, J., and Williams, K. C. W. (2025). IPUMS USA: Version 16.0 [dataset]. https://doi.org/10.18128/D010.V16.0.

Schulz, J., Agoha, C. E., Gebhard, A., Gregg, B., and Mayerhoffer, D. M. (2026). Excessive white male privilege biases the measurement of intersectional wage discrimination. *Humanities and Social Sciences Communications*, 13:631. https://doi.org/10.1057/s41599-026-06822-y.

Shorrocks, A. F. (1984). Inequality decomposition by population subgroups. *Econometrica*, 52(6):1369–1385. https://doi.org/10.2307/1913511.

Sprengholz, M. and Hamjediers, M. (2024). Intersections and commonalities: Using matching to decompose wage gaps by gender and nativity in germany. *Work and Occupations*, 51(2):249–286. https://doi.org/10.1177/07308884221141100.

Tomaskovic-Devey, D. (1993). The gender and race composition of jobs and the male/female, white/black pay gaps. *Social Forces*, 72(1):45–76. https://doi.org/10.1093/sf/72.1.45.

U.S. Bureau of Labor Statistics (2026). Consumer price index for all urban consumers (CPI-U), all items in u.s. city average. Bureau of Labor Statistics data series.

Vo, T., Schleifer, C., and Hekmatpour, P. (2023). Asian americans and income inequality: Disparities between and within racial, ethnic, and gender groups. *Sociological Perspectives*, 66(6):1103–1124. https://doi.org/10.1177/07311214231177019.