Perceptions of Race in the Labor Market*

Pedro C. Sant'Anna[†] Sulin Sardoschau[‡] Aiko Schmeißer[§]

May 2026

Abstract

Empirical studies of racial disparities often treat race as fixed, though scholars across disciplines view it as fluid, context-dependent, and observer-specific. We study the determinants of racial perception in the labor market and its consequences for measured wage disparities. Using linked administrative data on 330,000 workers in Brazil (2003–2015), we observe three racial measures: employer-recorded race from payroll filings, self-identified race from electoral records, and phenotype from standardized photographs. Self-identified and employer-ascribed race diverge in over 20% of cases, and employers disagree in their classification of the same worker. Estimating a "race function" that captures how employers map worker characteristics into racial categories, we find that phenotype matters, but so do productivity-relevant characteristics. University-educated workers are 10 pp more likely to be classified as White than workers with the same skin tone but no high school degree. Education "whitens" workers even conditional on self-declared race, all visual cues in the photographs, and firm-by-occupation fixed effects. Conventional wage-gap estimates based on recorded race or skin tone shrink by 80–90% under a comprehensive set of controls, partly because those controls also shape racial classification. We introduce a perception-normalized disparity measure that attenuates far less, suggesting conventional estimates overstate how much productivity differences explain racial wage gaps.

Keywords: Race, identity, disparity, wage gap, Brazil
JEL Codes: J15; J50; J71; Z10

*We would like to thank Christian Dustmann, Amy Finkelstein, Michael França, Larry Katz, Ben Olken, Frank Schilbach, as well as participants of the Labor Workshop at the 2025 RIDGE Forum, ifo Junior Labor Economist Workshop, 40th Meeting of the European Economic Association, 5th Arne Ryde Workshop on Culture, Institutions, and Development at Lund University, 1st Workshop on the Political Economy of Identity at RFBerlin, Columbia-Cornell Political Economy of Work Junior Scholars Workshop, MIT's Behavioral Lunch and Algorithms Coffee, Columbia's Labor & Public Finance Colloquium, and Research Seminar at the Immigration Policy Lab ETH Zurich. We also thank the George & Obie Shultz Fund at MIT.

[†]MIT Department of Economics and Núcleo de Estudos Raciais (INSPER). E-mail: p_stanna@mit.edu
[‡]Humboldt Universität zu Berlin, RFBerlin and IZA. E-mail: sulin.sardoschau@hu-berlin.de
[§]Columbia University, Columbia Center for Political Economy. E-mail: as8091@columbia.edu

1 Introduction

Racial disparities are pervasive across labor markets, education, criminal justice, and other domains, motivating a large empirical literature on discrimination. Most studies treat race as a fixed individual attribute. However, scholars across the social sciences have long viewed race as fluid and context-dependent. Racial classification can change over time and across settings, diverge between observers, and draw on cues that go well beyond phenotype or ancestry, including education, occupation, speech, and dress (Rose, 2023). This raises two central questions for research seeking to identify discrimination based on the race perceived by a decision-maker: What shapes racial perceptions? And what do disparity estimates capture when race is itself endogenous to social and economic status?

Addressing these questions empirically requires data that rarely exist: multiple measures of race for the same individual, including the classification held by the decision-maker, together with objective measures of appearance and socioeconomic status. Self-reports, the most frequently used measure in disparity estimates, capture how individuals classify themselves. Enumerator classifications and audit-study treatments capture how researchers or third parties classify them. None directly measures the perception held by the employer, evaluator, or judge whose behavior is under study. Without that measure, both the formation of racial perceptions and their implications for estimated disparities remain ambiguous.

This paper studies the determinants of racial perception in the labor market and its consequences for measured earnings gaps. We assemble linked administrative data for 330,000 workers in Brazil between 2003 and 2015, combining three distinct racial measures for each individual: employer-recorded race from mandatory payroll filings, self-identified race from electoral records, and algorithmic measures of skin tone and other phenotypic traits extracted from standardized photographs. These records are further linked to detailed information on wages, education, occupations, employment histories, and firms.

Our paper makes four contributions. First, we develop a conceptual framework that clarifies the role of worker self-identification and employer perceptions in wage-setting and shows how endogenous classification affects the measurement of disparities. Second, we document substantial divergence between how workers describe themselves and how employers classify them, as well as variation across employers in the category assigned to the same worker. Third, we estimate a race function that characterizes how employers map phenotype, self-identification, local context, and social status into racial categories. Fourth, we compare approaches to estimating wage gaps based on alternative measures of race and show that accounting for racial perceptions affects both the magnitude and interpretation of estimated disparities.

Brazil's fluid system of racial classification, shaped by its colonial history and socioeconomic structures (Telles, 2014), provides an ideal setting for examining how racial identity is constructed and interpreted across different institutional contexts, particularly in the labor market. At the same time, Brazil's experience reflects broader global patterns in the malleability of racial boundaries (Rose, 2023; Davenport, 2020). Shifts in self-identification in response to economic conditions, policy incentives, and evolving social norms appear in settings such as the United States (Baron et al., 2026; Adukia et al., 2025; Antman and Duncan, 2023; Dahis et al., 2019; Antman and Duncan, 2015), India (Atkin et al., 2021; Cassan, 2015), China (Jia and Persson, 2021), Indonesia (Rademakers and van Hoorn, 2021), and Romania (Mitrut et al., 2025).

We begin by developing a simple framework that makes the process of racial perception explicit and clarifies how it shapes measured disparities, building on an emerging literature on constructivism in economics (Rose, 2023). The framework distinguishes self-identification from employer-ascribed race and models each as the outcome of a classification process that maps phenotypic and status cues into racial categories. In some instances, these measures coincide, and in others, they diverge in systematic ways that make one measure more relevant than the other. The framework highlights that divergence may not simply reflect noise in reporting but captures meaningful differences in information sets or how workers and employers internalize or interpret the same cues.

The framework also clarifies why standard estimates of racial wage gaps are difficult to interpret. The implicit estimand targeted by the discrimination literature is the effect of being perceived as non-White by the employer, holding productivity fixed.[1] The standard conditional regression estimates this object by using observable characteristics as proxies for productivity. The difficulty is that productivity-relevant characteristics play two roles simultaneously: they affect wages through legitimate channels, and they shape how workers are racialized in the first place. Conditioning on education or occupation to hold productivity fixed therefore also holds fixed an input to racial classification, partialling out parts of race itself. In the limit, if these factors fully determine racial perception, the residual variation is pure noise. The meaning of the surviving race coefficient consequently hinges on a model of how racial perceptions are formed and does not provide unambiguous evidence that productivity differences alone explain the reduction in racial wage gaps.

Reduced-form and experimental designs that exploit a single cue, such as a distinctively Black name or skin tone, face a related challenge: they identify an intent-to-treat effect of the cue. The reduced-form effect compounds two distinct margins: how strongly the cue shifts the decision-maker's racial perception, and how strongly perceived race shifts the outcome. A null estimate is consistent with either margin being small, but only in the second case can this be read as evidence of no discrimination. Adding controls or providing additional background information on workers can shift the mapping from cue to perceived race. A smaller reduced-form effect after conditioning may therefore reflect either genuinely smaller disparities among similarly productive workers or a weaker first stage in which perceptions respond less to the cue once productivity-relevant characteristics are fixed. In either case, modeling the first stage – how racial perceptions are formed – is central to our understanding of the scope of these issues and their potential solutions.

In the remainder of the paper, we take this constructivist perspective on race to the data. We exploit information from *Relação Anual de Informações Sociais* (RAIS), which provides matched employer-employee records for the universe of formal jobs in Brazil with detailed information on workers' employment histories, firm characteristics and, crucially, includes the worker's race as reported by the employer. We merge RAIS to data on the universe of political candidates from the *Tribunal Superior Eleitoral* (TSE), for whom we observe self-

[1] This object aligns with the legal and economic notion of disparate treatment, which concerns differential treatment by a decision-maker on the basis of a protected characteristic as perceived at the time of the decision.

declared race.[2] Relative to the full RAIS workforce, this linked sample is slightly older and more male but comparable in education and average earnings. Labor market attachment remains strong because most candidates run for municipal council positions that require no more than 10 hours per week, and roughly 85% of candidates do not win office. Importantly, the TSE contains standardized photographs from which we extract an algorithmic measure of skin tone following the pipeline developed by Adukia et al. (2023). We also develop our own machine-learning prediction model that takes into account all the visual cues contained in the photos, such as workers' facial features.

We begin by documenting mismatched racial identities in the labor market. We show that self-declared and employer-recorded race are related but not interchangeable race concepts. About one quarter of workers are classified in different broad categories (White or non-White), and mismatches are asymmetric: employers are much more likely to record self-identified non-White workers as White than to classify self-identified Whites as non-White. These mismatches concentrate among individuals with intermediate skin tones and are closely tied to job mobility, with a sizable share of workers changing racial classification when they move between firms.[3] In about 20 percent of cases, employers do not agree on the racial category of the same worker. We provide evidence that measurement error alone is unlikely to explain this mismatch.

To examine how employers ascribe race, we estimate a "race function" that maps workers' skin tone, self-identification, and a broad set of demographic and labor market characteristics into recorded race. Controls include age, gender, education, past wages, and tenure, along with region, industry, occupation, and firm fixed effects. First, we show that phenotype matters: the probability of being classified as non-White rises monotonically with skin tone, and a one standard deviation increase raises the likelihood of non-White ascription by 11 percentage points. Using the full set of visual features extracted from photographs via machine learning triples the explanatory power of photo-based cues. Second, consistent with the patterns of mismatched identities documented above, self-identification and employer ascription are distinct objects. Self-identification strongly predicts recorded race, but employers systematically depart from it: they are more likely to disagree with workers who self-identify as White when those workers have a darker skin tone.

Importantly, and consistent with endogenous racial perception, productivity-relevant characteristics shape employer-recorded race. Conditional on skin tone and the full set of visual features, education, prior earnings, and labor market tenure are associated with the

[2] Ideally, we would measure the ascribed race by the actual wage-setter at the moment of the pay decision, and self-declared race as the racial identity the worker presents in public economic life. Our proxies come close to these constructs: employer-ascribed race in RAIS is typically recorded by the firm owner or manager at the moment of hiring, and self-declared race in TSE captures workers' racial self-presentation in formal public life. However, they are not perfect: RAIS reflects a firm-level rather than wage-setter-specific perception, while TSE self-declaration is made under candidacy-specific stakes and in a selected sample of workers who run for office. We discuss these issues in Sections 3.2, 5 and 6.3, and Appendix C, and probe them empirically where possible.

[3] To quantify the role of firms in racial ascription, we extend the analysis to the full RAIS sample of roughly 45 million workers and exploit worker mobility across establishments in an AKM framework (Gerard et al., 2021; Abowd et al., 1999). Modeling employer-recorded race with worker and firm fixed effects shows that establishment-specific classification practices account for about one-third of the variation in recorded race, a contribution that would be absent if race reflected an innate worker attribute.

"whitening" of workers. For example, university-educated workers are 10 percentage points more likely to be classified as White than workers without a high school degree of the same skin tone. These patterns persist within firm-by-occupation cells and conditional on self-declared race, suggesting that racial perception is formed at the job level, in the interaction between worker and employer, rather than through sorting of workers across occupations and firms.

In the final part of the paper, we explore how racial perception affects the measurement of wage disparities. We begin by estimating conventional wage gaps in RAIS using employer-recorded race, successively introducing the same controls and fixed effects as in the race-function analysis. Workers classified as non-White by their employer receive 14.8% lower wages on average. The disparity shrinks by almost 90% in the most stringent specification, with a remaining 2.1% wage gap between ascribed non-Whites and Whites who have the same demographic characteristics and work in the same firm and occupation. Finally, we augment the wage disparity specification with self-declared race and find that conditional on self-declared race, employer-ascribed race carries an additional wage penalty of 1.7%, emphasizing that employer classification is an important margin in wage setting. Importantly, and as discussed in our conceptual framework, the large reduction of wage gaps when adding controls can be explained by two forces: *i*) differences in productivity or sorting of workers and *ii*) controls netting out determinants of racial classification itself. Thus, the interpretation of conditional disparity estimates depends crucially on the nature of the remaining variation in racial perceptions.

We then turn to an alternative approach that is more explicit about the exploited variation, focusing on one specific racial cue. Following the colorism literature, we study wage disparities along the skin tone distribution, using our algorithmic measure from standardized photographs. We document a near-monotone relation: a one standard deviation darker skin tone is associated with a 2.6% wage penalty. Similar to disparities in terms of employer-recorded race, the skin tone disparity shrinks by about 80% in the most stringent specification. However, as noted in the conceptual framework, cue-based disparities cannot distinguish weaker racialization from weaker discrimination. In fact, we document the former: the coefficient of skin tone on employer-ascribed race attenuates by about 40% relative to the unconditional estimate. Most notably, the drop in skin tone disparity coincides with the drop in the skin tone elasticity of the race function. This underscores that cue-based disparities should be interpreted as an intent-to-treat effect whose impact depends on how strongly the racial cue translates into employer perceptions.

Building on these insights, we estimate a perception-normalized disparity (PND) that scales the reduced-form impact of skin tone on wages by the corresponding first-stage relation between skin tone and employer's race perception. Conceptually, the PND expresses disparities in units of racialization rather than in units of physical appearance. The unconditional PND reveals a 23.5% wage penalty compared to 14.8% when using employer-ascription directly. In our most stringent specification that includes firm-by-occupation fixed effects and demographics, we estimate a PND of 8.0% which is only about 65% smaller than the unconditional PND. As expected, this attenuation is considerably smaller than for the conventional measures, in line with the idea that the PND estimates restore perception-driven variation among workers with otherwise similar characteristics. These results highlight how conventional approaches can overstate the role of differences in productivity (or worker and job characteristics more generally) in explaining racial wage gaps.

Two further analyses underscore the value of accounting for racial perception. First, racial cues are interpreted heterogeneously: the elasticity of skin tone in the race function is 15% smaller for university-educated workers than for those without a high school degree. This means that employer-recorded race is less sensitive to skin tone for workers of higher social status. Consequently, the colorism-disparity is smaller than the corresponding PND across these groups. More generally, the PND allows for meaningful comparisons of disparity across workers with different characteristics and employers with different race functions. Second, a split-sample IV using repeated photographs more than doubles the reduced-form effect of skin tone on wages, revealing substantial attenuation bias in colorism disparity due to measurement error, something that the scaling approach can absorb.

We also discuss the PND's limitations. Interpreting the PND as the causal effect of racial perceptions requires an exclusion restriction: the cue must affect wages only through perceived race. This may fail if observationally similar workers within firm-by-occupation cells still differ in unobserved productivity, or if employers hold direct preferences over skin tone independent of racial categories. Even under a valid exclusion restriction, the PND recovers a cue-specific local effect that need not generalize to names or other signals.

Overall, our findings suggest that identity is not only something individuals bring to markets; it is partly something that markets produce. This has implications for two of the most studied objects in labor economics: returns to education and the measurement of discrimination. Returns to education may include not only productivity gains and signals of ability but also returns to racial perception. The measurement of discrimination is affected by the same mechanism: controlling for education is not neutral, since education is itself one of the cues that shape how race is perceived. The same logic applies wherever people judge others by combining multiple signals into a single category, whether race, religion, caste, or ethnicity, and extends to settings beyond the labor market, including criminal justice, credit, and healthcare.

Our paper contributes to the literature on the endogeneity and fluidity of racial classifications, but speaks to broader research on the economic consequences of endogenous social identity (Oh, 2023; Del Carpio and Guadalupe, 2022; Lowe, 2021; Shayo, 2020; Atkin, 2016; Ben-Ner et al., 2009; Akerlof and Kranton, 2000). On the theory side, a growing line of research argues that, because race is socially constructed, the well-defined causal object in discrimination research is the perception of race at a specific moment of decision rather than race as a fixed individual attribute (Hu and Kohler-Hausmann, 2025; Rose, 2023; Kohler-Hausmann, 2019; Sen and Wasow, 2016; Greiner and Rubin, 2011; Holland, 2008). On the empirical side, many papers document individual race changes over time and associate these changes with incentives, status, and context, for instance, in the United States (Davenport et al., 2026; Adukia et al., 2025; Noghanibehambari and Fletcher, 2025; Antman and Duncan, 2023; Davenport, 2020; Dahis et al., 2019; Antman and Duncan, 2015; Saperstein, 2012; Saperstein and Penner, 2012; Penner and Saperstein, 2008) and in Brazil (Freeman et al., 2025; Muniz and Bailey, 2022; De Micheli, 2021; Chor et al., 2019; Cornwell et al., 2017; Miranda, 2015; Mitchell-Walthour and Darity Jr., 2014; Telles and Paschel, 2014; Francis and Tannuri-Pianto, 2013; Schwartzman, 2007). The empirical literature mostly focuses on changes in self-declared race, and previous studies of hetero-identification were restricted to laboratory settings or captured only enumerators' racial perceptions and do not contain objective phenotypic features (Telles and Lim, 1998; Bailey et al., 2013; Rodeheffer et al., 2012; Freeman et al., 2011; Saperstein and Penner, 2012). In contrast, by jointly observing self-identification, algorithmically extracted phenotype, and the race recorded by a relevant decision-maker, we provide the first large-scale empirical operationalization of constructivism, providing at least two key contributions to this literature. First, we document that employers' racial perceptions systematically depart from self-declared race and phenotype, reflecting in part social status cues. Second, we show that the choice of racial measure matters quantitatively for estimated disparities.

Our paper also adds to studies on the measurement and interpretation of racial disparities in labor markets and other settings. One line of research uses observational data to decompose racial gaps in wages and employment by regressing labor market outcomes on a race indicator while controlling for proxies of productivity (Bayer et al., 2025; Sorkin, 2025; Bayer and Charles, 2018; Card et al., 2018; Fryer et al., 2013; Fryer, 2011; Darity Jr et al., 1996; Blinder, 1973; Oaxaca, 1973). A parallel experimental literature identifies discrimination from audit and correspondence studies that manipulate racial cues such as names, photos, or speech in applications or interactions and measure differential responses by employers (Ajzenman et al., 2025; Evsyukova et al., 2025; Kline et al., 2022; Dias, 2020; Agan and Starr, 2018; Neumark, 2018; Quillian et al., 2017; Doleac and Stein, 2013; Pager, 2007; Bertrand and Mullainathan, 2004). A related theoretical literature provides frameworks for interpreting what these empirical estimates capture (Bohren et al., 2025a; Arnold et al., 2022; Small and Pager, 2020; Bohren et al., 2019; Arnold et al., 2018; Aigner and Cain, 1977; Arrow, 1973; Phelps, 1972; Becker, 1957). We contribute to this literature by showing that racial perceptions matter for how we can interpret both observational decompositions and audit-study estimates: benchmarking estimates shift systematically depending on how controls shape racial perceptions and audit cues identify the effect of a signal rather than of perceived race itself. Our results suggest that conventional approaches may overstate the role of productivity proxies and highlight the importance of measuring racial perceptions.

Finally, our work also contributes to the literature on the measurement and interpretation of racial signals. One line of research, known as colorism, estimates the effect of skin color on labor-market and other outcomes. These studies typically use observer-reported complexion and document that darker skin is associated with worse outcomes within and across racial groups in both Latin America (Woo-Mora, 2026; Monk, 2016; Telles, 2014, 2004) and the US (Adukia et al., 2025; Kreisman and Rangel, 2015; Monk, 2014; Hersch, 2008, 2006; Goldsmith et al., 2007, 2006).[4] A related line of research seeks to isolate the effects of other racial cues such as names or hairstyle (Kreisman and Smith, 2023; Koval and Rosette, 2021; Fryer Jr and Levitt, 2004). We contribute to this literature in two ways. First, by using algorithmically

[4]A paper that goes beyond estimating colorism, and is closer to ours, is Adukia et al. (2025). Studying the social construction of race in a historical US setting, they link bank tellers' written descriptions of depositors' complexion to enumerator-reported race in the 1870 Census and ask whether wealth differences correlate with census-recorded race conditional on recorded skin tone. We extend this work in three respects. First, our setting is a contemporary labor market. Second, rich administrative data let us measure phenotype algorithmically, rather than through a second observer, alongside racial self- and outside-perceptions, detailed job and firm characteristics, as well as earnings trajectories. Third, we observe the perception of a consequential economic decision-maker, the employer, and we measure outcomes directly tied to that decision environment.

extracted skin tone, we break the circularity in which observer-reported phenotype is itself shaped by the perceptual processes the literature aims to study. Second, we show that racial signals are not interpreted uniformly across observers and subjects: they vary in how strongly they shift perceptions, and that strength can itself depend on the subject's other characteristics.

The remainder of the paper is structured as follows. Section 2 develops the conceptual framework. Section 3 provides background on race in Brazil and describes the data and construction of our racial identity measures. Section 4 documents mismatches between self-declared and employer-ascribed race and across employers, and Section 5 estimates the employer "race function." Lastly, Section 6 analyzes the implications for measuring racial wage disparities, and Section 7 concludes.

2 Conceptual Framework

We start by introducing a constructivist approach to race in the labor market. Building on the framework proposed by Rose (2023), we distinguish between self-identified race and employer-ascribed race and analyze when and why these measures coincide or diverge. We show that when racial perceptions depend on social cues, the interpretation of conventional disparity measures may change, as variables that directly affect outcomes may also shape the process of racial categorization.

Self-Identification and Employer-Perceptions. Let R^g denote the race category recorded by source $g \in \{S, E\}$, where S indexes self-identification and E indicates the race ascribed by the decision maker – in our case, the employer. We consider a binary coding

$$R^g \in \{0, 1\},$$

where 0 corresponds to White and 1 to non-White.[5]

We model racial classification as the discretization of a latent racialization index. Specifically, observer g maps a set of racial cues $Z^g \in \mathcal{Z}_g$ into a scalar index

$$a^g \;=\; A_g(Z^g) + u^g \in \mathbb{R},$$

where $A_g : \mathcal{Z}_g \to \mathbb{R}$ is the mapping of cues into a latent race index, and where u^g captures noise in the form of administrative inattention or recording errors. The recorded category R^g is then generated by a threshold rule,

$$R^g \;=\; \begin{cases} 1 & \text{if } a^g \geq \tau_g, \\ 0 & \text{if } a^g < \tau_g, \end{cases}$$

with $\tau_g \in \mathbb{R}$ an observer-specific threshold that captures category boundaries in a given social and institutional environment. In this formulation, Z^g may include phenotypic features (e.g., skin tone or hair texture), ancestry (e.g., family origin or surname-based cues), and status

[5] In principle, the model can be extended to accommodate multiple racial categories. In Section 3, we justify the use of a binary racial categorization in the Brazilian context.

signals (e.g., education, income, dress or speech). Moreover, the interpretation and weighting of these cues and their aggregation into racial categories, through $A_g(\cdot)$ and τ_g, may be shaped by contextual factors in which these interactions take place (e.g., location or firm).

This setup makes clear when self- and employer-perceptions can be treated as equivalent. Suppose there exists a common information set Z and a common mapping (A, τ) such that both sources implement the same classification rule,

$$R^S = \mathbf{1}\{A(Z) + u^S \geq \tau\}, \qquad R^E = \mathbf{1}\{A(Z) + u^E \geq \tau\},$$

then R^S and R^E differ only because of observer-specific noise. In the limiting case without such noise, $R^S = R^E$ and the two racial concepts are identical. This equivalence may hold in settings where racial cues are entirely unambiguous, or where individuals internalize repeated external feedback so that self-identification converges to dominant outside perceptions.

In other settings, however, equivalence need not hold. Self-identification and employer ascription can differ because the underlying objects determining R^g differ across observers. First, information sets can differ ($Z^E \neq Z^S$). Employers typically rely on observable cues in the hiring or wage-setting environment, whereas self-identification may incorporate private information about ancestry, family history, or social affiliation. Second, even when the same cues are available, observers may weight them differently or apply different category boundaries. This corresponds to differences in $A_S(\cdot)$ versus $A_E(\cdot)$ and in thresholds τ_S versus τ_E. In all these cases, disagreement between R^S and R^E is not merely measurement error around a single underlying "true" race. Rather it reflects heterogeneity in which cues are available and how they are translated into racial categories.

Which measure is most relevant depends on the research question. If the objective is to describe broad labor market inequality associated with racial identity, and one believes that self-identification captures channels such as differential social networks or behavioral responses to anticipated or past discrimination, then R^S may be the more relevant margin. If instead the objective is to study disparate treatment by decision-makers, then employer-ascribed race R^E is the relevant object because it is the category that enters the decision environment. In either case, choosing between R^S and R^E is a substantive modeling choice about the perception process, not a purely technical decision about measurement.

Conventional Disparity Measures. Conventional measures of disparity compare mean outcomes across racial categories, either unconditionally or conditional on other individual characteristics. If the goal is to describe differences in outcomes across recorded racial categories then these quantities are well-defined descriptive objects.

If the goal is, however, to identify *disparate impact* in unconditional estimates (capturing both structural/institutional and direct forms of discrimination) and *disparate treatment* in conditional estimates (direct discrimination at a specific decision-node), then researchers must take into account the process of racial categorization and correlations between recorded race and unobserved factors that may simultaneously affect the outcome of interest.

More specifically, researchers attempting to estimate the "causal effect of race" face a longstanding disagreement over what it means to estimate such an effect (e.g., Hu and Kohler-Hausmann, 2025; Rose, 2023). There are at least two reasons to hesitate to interpret recorded race itself as a cause. First, race is fluid, observer-specific, and context-dependent. Second, the notion of a counterfactual race is conceptually difficult. For these reasons, researchers often prefer to interpret the causal effect of race as the causal effect of *racial perceptions*, particularly when studying direct discrimination. This approach asks what would happen if a relevant decision-maker perceived the same individual differently. In this section, we focus on employers as they are the relevant decision-makers in our labor market setting, but the key insights are equally applicable to self-perceptions of race.

Let $Y \in \mathbb{R}$ denote a labor market outcome of interest, such as (log) wages or an indicator for employment. X denotes productivity-relevant characteristics of the worker. Then, the conditional racial disparity based on employer perceptions is defined as

$$\Delta(x) \;=\; \mathbb{E}[Y \mid R^E = 1, X = x] \;-\; \mathbb{E}[Y \mid R^E = 0, X = x].$$

The constructivist approach clarifies that racial disparity measures depend on how R^E is generated, because R^E is endogenous to observable cues and contextual information. Specifically, the key empirical complication is that the cue set Z^E can contain characteristics that are also productivity-relevant, i.e., $Z^E \cap X \neq \varnothing$. This creates an identification problem in conventional disparity estimates. If the target is the effect of employer-ascribed race holding productivity-relevant characteristics fixed, omitting elements of Z^E that affect outcomes and are correlated with R^E confounds perceived race with omitted productivity-relevant variation.

If, instead, the researcher observes and controls for these elements, the omitted-variable concern is reduced, but the interpretation of this conditional comparison changes: the comparison is conducted while holding fixed the inputs that generate the perceived racial category in the first place. When R^E is highly predictable from X, conditioning on X leaves little identifying variation in R^E. As controls absorb inputs into racial classification, the comparison increasingly holds fixed the cues through which employer-ascribed race is generated, and the interpretation would depend on what remaining variation in racial perceptions exists.

To illustrate this issue, let us take the example of education. Suppose schooling increases productivity, thereby directly affecting wages. At the same time, suppose lower schooling shifts employer classification toward non-White by entering $A_E(Z)$ as a salient status cue. If schooling is unobserved by the researcher, the estimated gap between $\mathbb{E}[Y \mid R^E = r^1]$ and $\mathbb{E}[Y \mid R^E = r^0]$ loads both on differences in perceived race and on differences in schooling that affect productivity.[6] If schooling is observed and included as a control, the confounding channel is addressed, but the estimand now compares workers who are similar in a key input into racial classification. In settings where schooling strongly explains R^E, conditioning on schooling removes much of the variation that generates the perceived racial identity. In the limiting case where employer classification is based entirely on schooling (plus random noise), conditioning on schooling eliminates all disparity by R^E, not because employers ignore race, but because schooling fully explains employers' racial perceptions of workers. More generally, when controlling for productivity-relevant factors, it is crucial to consider

[6] Note that, from a constructivist perspective, confounding arises because education influences racial perceptions and wages. This is different from the argument for controlling for education to identify disparate treatment by an employer when holding productivity differences constant. In fact, the endogeneity of racial perceptions to education also implies biased estimates of disparate impact when examining unconditional disparities.

what the remaining variation in racial perceptions (beyond noise) is and how it changes the interpretation of identified disparities.

Cue-Based Disparity Measures. An alternative approach that sidesteps direct race measures and is more explicit about the variation exploited, focuses on a single salient cue $z \in Z$. Colorism and audit studies, for example, study disparities by skin tone and names, respectively. Let us simplify to a case where $z \in \{z^0, z^1\}$ denotes two values of a cue. This strategy – which we call cue-based disparity – compares outcomes across values of z, often conditional on X:

$$\Delta_z(x) \;=\; \mathbb{E}\big[Y \mid z = z^1, X = x\big] \;-\; \mathbb{E}\big[Y \mid z = z^0, X = x\big] \,.$$

The interpretive challenge is that $\Delta_z(x)$ is not, by itself, a disparity in perceived race. If z is as-good-as random conditional on X, $\Delta_z(x)$ can be interpreted as an intent-to-treat (ITT) effect of the cue. Thus, it is informative about perception-based disparities only to the extent that the cue shifts employer-ascribed race. The sensitivity of racial ascription to the cue can be summarized by the first-stage, or what we call the "race function:"

$$\pi_z(x) \;=\; \mathbb{E}\big[\mathbf{1}\{R^E = 1\} \mid z = z^1, X = x\big] \;-\; \mathbb{E}\big[\mathbf{1}\{R^E = 1\} \mid z = z^0, X = x\big] \,.$$

When $\pi_z(x)$ is small, the cue rarely changes perceived race in that group, so even substantial disparate treatment based on R^E can generate only a small reduced-form disparity in z. Conversely, when $\pi_z(x)$ is large, the same cue difference induces larger changes in perceived identity, making cue-based disparities more informative about perception-based disparities. For these reasons, cue-based estimates are best interpreted jointly with evidence on how the cue maps into perceived race.

A simple thought experiment illustrates why $\Delta_z(x)$ and $\pi_z(x)$ should be considered jointly. Consider skin tone as the cue z, and focus on a group defined by $X = x$. The reduced-form skin tone disparity can be viewed as the product of two components: (i) how strongly skin tone shifts racial ascription, captured by $\pi_z(x)$, and (ii) how strongly wages respond to being perceived as non-White. A near-zero reduced-form skin tone disparity $\Delta_z(x) \approx 0$ can therefore arise in two distinct states of the world. In a de-racialized or "color-blind" world, skin tone does not predict racial categorization, so $\pi_z(x) \approx 0$ and racial categories have been effectively dissolved. In a racialized but non-discriminatory world, skin tone strongly predicts racial categorization, so $\pi_z(x)$ is large and racial categories are maintained, but wage-setting does not differ by perceived race, so the disparity conditional on perceived identity is near zero. The reduced-form estimates in z alone cannot distinguish these cases.

In addition, $\pi_z(x)$ can vary systematically with X. This means that the same racial cue can carry a different weight depending on the regional context or other worker characteristics, such as education or income. For example, a darker skin tone may shift employer classification toward non-White more strongly among low-educated workers than among highly educated workers. If the estimated reduced-form effect of skin tone on wages is close to zero among highly educated workers, this may reflect genuinely smaller disparities conditional on being perceived as non-White. But it may also reflect that skin tone is less likely to change perceived race in the first place among the highly educated.

Perception Normalization. This suggests a simple normalization that puts cue-based gradients on a common scale. For a binary cue $z \in \{z^0, z^1\}$, define the *Perception-Normalized Disparity* (PND) as

$$\mathrm{PND}_z(x) \;\equiv\; \frac{\Delta_z(x)}{\pi_z(x)} \;=\; \frac{\mathbb{E}[Y \mid z = z^1, X = x] - \mathbb{E}[Y \mid z = z^0, X = x]}{\mathbb{E}[\mathbf{1}\{R^E = 1\} \mid z = z^1, X = x] - \mathbb{E}[\mathbf{1}\{R^E = 1\} \mid z = z^0, X = x]}.$$

Intuitively, $\mathrm{PND}_z(x)$ rescales the outcome gradient in the cue by the extent to which the cue changes the likelihood of being perceived as non-White. When $\pi_z(x)$ is close to zero, the cue does not move perceived race in that group, so cue-based disparities are not informative about perception-based disparities along that margin, and the normalization makes this lack of a racialization margin explicit. Compared to the reduced form, the $\mathrm{PND}_z(x)$ expresses disparities in units of perception rather than units of the cue.

Importantly, $\mathrm{PND}_z(x)$ should be interpreted as a scaling exercise rather than the causal effect of perceived race unless one can argue that the cue affects outcomes only through employer racial ascription. Under such an exclusion restriction, $\mathrm{PND}_z(x)$ coincides with an instrumental-variables estimand and can be interpreted as a local average treatment effect (LATE) of being perceived as non-White for observations whose perceived race changes when the cue shifts. The estimand changes depending on the specific racial cue used and does not need to be the same, such that, for instance, $\mathrm{PND}_{z=\text{skin tone}}(x) \neq \mathrm{PND}_{z=\text{name}}(x)$.

Despite these caveats, the PND's scaling approach helps make cue-based disparity estimates more comparable across different specifications, sub-samples, and measures of racial cues. First, if the racial cue z is correlated with covariates X that also influence employers' racial perceptions, conditioning on X may alter the sensitivity of perceived race to the cue. As a result, a smaller reduced-form estimate in a conditional comparison may not only reflect differences in productivity-relevant characteristics across groups but also a weaker racialization margin. The PND incorporates the attenuated first stage, thereby making results more comparable across conditional contrasts. Second, the PND accommodates heterogeneity in how strongly the cue shifts perceived race across workers with different X (such as in the education example above). Third, noise in the measurement of the racial cue attenuates both the first-stage and the reduced-form estimates, while leaving the PND unaffected.

In sum, our conceptual framework suggests that conventional disparity measures based on R^g are inherently sensitive to control choices whenever cues that enter the racial classification rule are also productivity-relevant ($Z \cap X \neq \varnothing$). Cue-based disparities $\Delta_z(x)$ are, in turn, sensitive to the strength with which a given cue shifts perceived race, which may vary systematically across worker characteristics and contexts. Abstracting from data constraints, each of these measures is well-defined for a distinct estimand and rests on different identifying assumptions. Our empirical analysis does not seek to rank these measures, but to quantify the magnitude of the resulting measurement differences in a real-world setting and to make explicit how assumptions about racial perception shape both the interpretation and the stability of conventional disparity estimates.

3 Background and Data

3.1 Race in Brazil

Race in Brazil is shaped by the country's colonial history and institutional legacy. Brazil's racial system originated during the colonial era, when Portuguese colonizers enslaved millions of Africans and forced Indigenous peoples into labor. The country was the last in the Americas to abolish slavery in 1888, but formal segregation policies after abolition never matched those of the United States. Yet, the absence of formal segregation did not eliminate racial hierarchy. Influential scholars such as Freyre (1933) advanced the idea of "racial democracy," suggesting that widespread miscegenation fostered more harmonious race relations than in the United States. Critics like Schwarcz (2013), Fernandes (1965) and Holanda (1936) highlighted persistent inequalities and argue that Brazil's so-called racial democracy obscured socioeconomic gaps separating White and non-White populations.

In official statistics, the Brazilian Institute of Geography and Statistics (IBGE) recognizes five racial categories: *branco* (White), *preto* (Black), *pardo* (Brown), *amarelo* (Asian), and *indígena* (Indigenous). However, racial boundaries in Brazil typically follow a White versus non-White divide (pooling *preto* and *pardo*). Asian and Indigenous individuals make up only about one percent of the Brazilian population (IBGE, 2022). Sociological and ethnographic work emphasizes that discrimination is organized around the boundary of whiteness, whereas distinctions within Afro-descendant categories are comparatively narrow and context-dependent (Schwarcz, 2013; Telles, 2004; Fernandes, 1965). This aggregation is also conventional in the literature quantifying racial disparities in Brazil (e.g., Miller and Schmutte, 2023; Gerard et al., 2021; Firpo et al., 2021). Policy design mirrors this practice: affirmative-action rules typically treat *pretos* and *pardos* jointly as beneficiaries, rather than targeting them separately (Francis and Tannuri-Pianto, 2013). Moreover, it is reflected in measured inequalities. *Preto* and *pardo* workers are shown to exhibit relatively similar levels of education, employment, and wages (e.g., Miller and Schmutte, 2023; Gerard et al., 2021; Silva, 2000). At the same time, both groups experience substantially worse socioeconomic outcomes than White individuals: non-Whites earn 35% less, are more likely to be unemployed, are underrepresented in leadership positions, and have fewer years of schooling (França and Portella, 2023; IBGE, 2022).[7]

There is also mounting evidence that race in Brazil is not static. Individuals and observers classify or perceive the same person differently depending on economic conditions or institutional settings. In the last two decades, many Brazilians changed their self-identification towards *preto* and *pardo* (Miranda, 2015), possibly in response to increased opportunities and lower stigma (De Micheli, 2021) or policy incentives (Francis and Tannuri-Pianto, 2013). At the same time, improvements in socio-economic status lead self- and external-identification to move towards White (Cornwell et al., 2017; Schwartzman, 2007; Maggie and Fry, 2004).

[7] Appendix Table B.3 confirms that, in our sample of political candidates, socio-demographic and labor-market profiles are much more similar for *preto* and *pardo* individuals than for *branco* versus *pardo*.

3.2 Data

We combine information from two large-scale administrative data sources that allow us to observe three racial identity measures for the same individual. First, we draw on the *Relação Anual de Informações Sociais* (RAIS) that provides matched employer-employee data on all formal jobs in Brazil and, importantly, records workers' race *as reported by the employer*. Second, we merge RAIS with data on the universe of political candidates in Brazil, for whom we observe their *self-declared race* as well as algorithmic measures of *phenotype* derived from standardized electoral photographs. This section describes the two data sources, the three racial identity measures, and the steps used to construct our estimation sample.

Employer-Ascribed Race. We use RAIS records covering all formal sector employment relations in Brazil between 2003 and 2015. The data is collected annually by Brazil's Ministry of Labor and Employment for the administration of various social security programs. All employers are mandated to submit information on their employees via an online platform provided by the Ministry (Appendix Figure C.1 shows a screenshot of the interface). Compliance with reporting requirements is high because incomplete submissions can trigger substantial penalties (MTE, 2024).

Each record in the data corresponds to a worker-establishment contract in a given year and contains detailed information on average monthly earnings, contracted hours, occupation, contract type, the start and end date of the contract, along with workers' sociodemographic characteristics (gender, age, education) and the firm's industry and location. Workers are identified by their unique tax identifier *Código de Pessoa Física* (CPF) and their names, and can be tracked across employers.

Crucially, RAIS collects data on employees' race since 1999. The race question asked in RAIS and the available racial categories are the standard ones in Brazilian statistics, as discussed in the previous section. There is no specific guidance on how to report a worker's race, either in the official RAIS Manual or in the reporting interface, and the reported race information is not subject to any systematic auditing (MTE, 2024). Most employers do not seem to consult workers and instead classify race based on their own perceptions at the time of hiring (Silveira, 2022; Osório, 2003). See Appendix C.1 for more details on how employers collect and report worker information, including results from qualitative interviews that we conducted with employers. Reporting of RAIS is done at the establishment level, by a dedicated HR or administrative employee, or by the owner or manager themselves in the case of smaller establishments (Silveira, 2022; Osório, 2003). Thus, the race measure should be interpreted as an establishment-level report of worker race, rather than necessarily reflecting of the specific decision-maker responsible for wage setting within the firm.

There are no direct economic incentives for reporting a more or less racially diverse workforce (e.g., no tax incentives or affirmative action laws in the private labor market). Hence, we do not expect strategic reporting in the sense of changes in employees' ascribed race to achieve a desired racial composition. At the same time, the presence of a race question in RAIS is not unusual for respondents, as Brazilians have long been accustomed to being asked about their race and that of others in official statistics. Taken together, these features suggest that employers take reporting of race in RAIS seriously and that employer-reported race is a good proxy for how employers perceive the racial identity of their workers.

Nevertheless, we will subject our main results to a series of robustness checks designed to explicitly account for potential measurement error in employer-ascribed race.

Self-Declared Race. To enrich RAIS with independent measures of racial identity, we link it to administrative records from Brazil's *Tribunal Superior Eleitoral* (TSE), which contain information on the universe of individuals who run for political office. Beginning in 2014, the TSE registry includes each candidate's self-declared race using the five standard categories. Importantly, only about 15% of all candidates are elected, and about 90% run for municipal council positions. These are typically part-time roles – on average, municipal councilors are required to attend the council only four days per month (Ferraz and Finan, 2011) – during which councilors maintain their private-sector jobs. The strong labor market attachment makes the candidate sample well-suited for comparing self-declared and employer-ascribed racial identities. Candidates in the TSE records are identified by their tax identifier (CPF), enabling a deterministic linkage to RAIS and avoiding any potential mismatch in racial identities that might arise under probabilistic record-linkage procedures.

Candidate registration is an administrative declaration to the TSE, publicly available on the TSE website and electoral apps. Self-declared race in TSE records is therefore a relatively public signal, though candidates' race is not displayed on the ballot, and prior work finds limited electoral returns to self-declaring race in TSE registration (Janusz, 2023).[8] To the extent that the signal is public, employers may plausibly observe it. Self-identification is also not necessarily fixed, and racial self-classification can change over time (Janusz, 2021). In the empirical analysis, we examine whether strategic considerations or changes in self-identification help account for mismatches in racial identity.

Skin Tone. The candidate registry also provides standardized photographs used in electronic voting machines for nearly all candidates between 2004 and 2024, yielding close to three million photos. We use these head-shots to construct an algorithmic measure of skin tone that is derived directly from the color value of each pixel in a photo. In contrast to existing measures in the literature, which typically rely on enumerator-reported skin tone (Adukia et al., 2025; Bueno and Dunning, 2017; Telles, 2014), this approach delivers an objective and reproducible measurement of skin tone.[9] This circumvents the issue that characteristics and experiences of the enumerator or respondents' other racial cues affect how racial cues are reported in the first place.

The algorithmic pipeline is illustrated in Figure 1. We first isolate the skin region in each photo using Google MediaPipe's Image Segmentation tool, which segments the skin

[8] This is not to say that the racial categories candidates select are immune to public image concerns. We interpret self-declared race as the identity individuals choose in public life, shaped by both internalized identity and social desirability. A similar interpretation may apply to identities reported to enumerators in standard survey data (Pickett et al., 2023).

[9] There are few other papers using algorithmic procedures to extract some racial identity measure from TSE's photos (Baerlocher and Schneider, 2026; Pérez-Cervera, 2025; de Lucena Coelho et al., 2024). All these papers use off-the-shelf "race-detection" algorithms to predict broad racial categories, with the goal of either estimating the effects of electing non-White candidates or of using candidates' race as an outcome. In contrast, we extract specific phenotypic traits from photos or use a fully flexible machine learning prediction model to understand how racial perceptions form.

part from other elements of the image (background, hair, clothing, accessories). We then summarize the color of the skin region by closely following the algorithm of Adukia et al. (2023). We use a k-means clustering approach to group pixels into $k = 5$ clusters and exclude the smallest two clusters that often refer to shadows or other non-skin facial features. For each remaining cluster, we compute its characteristic color and express these values in the $L^*a^*b^*$ color space. A weighted average of these cluster colors produces a single representative skin color for each candidate. Appendix C.2 provides more information on the characteristics of electoral photos, details on the measurement procedure, as well as a validation exercise that compares our algorithmic skin tone measure to human assessments of candidates' skin tone collected by Bueno and Dunning (2017).[10]

Figure 2 depicts the distribution of skin tone in our data, separating self-declared White and non-White candidates (Appendix Figure C.2 also shows example photos across the skin tone distribution, and Appendix Figure A.2 splits the non-White group between *pardos* and *pretos*). Skin tone has a mean of 38.6 and a standard deviation of 10.8, with self-declared White and non-White candidates differing by an average of 0.62 standard deviations. Throughout the analysis, we standardize skin tone and report estimates in standard deviation units.

Figure 1: Skin Tone Extraction Pipeline

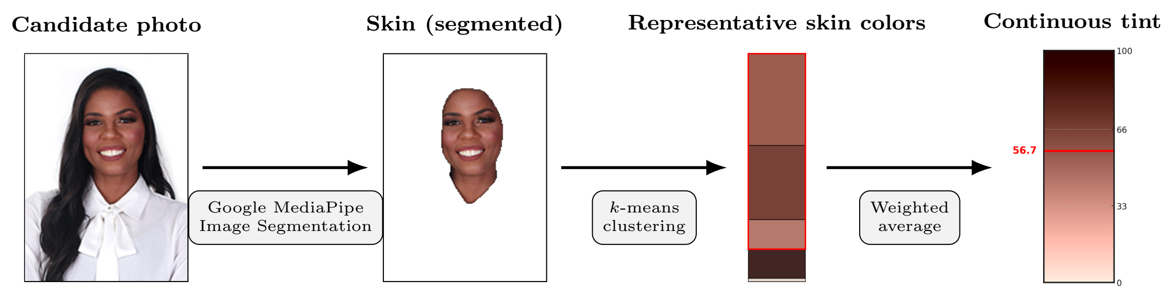

Notes: The figure illustrates how we extract skin tone from electoral photos. The example shown corresponds roughly to the 90th percentile of the skin tone distribution. Appendix Figure C.2 provides additional examples across the distribution.

Machine-Learning Prediction of Visual Cues. In addition to skin tone, we also construct a broader measure of racial cues that capture all information visible in the candidate photos. Specifically, we implement a machine-learning procedure that predicts employers' racial perceptions using the full set of facial features visible in the photos. The algorithm follows a transfer learning approach: it first extracts facial embeddings from each photo using a pretrained face-recognition encoder (Serengil and Ozpinar, 2024; Schroff et al., 2015), and then combines these embeddings with our skin tone measure to train an RBF-kernel SVM model (Fournier-Montgieux et al., 2025). Appendix C.3 provides details on the prediction procedure. Given the black-box nature of the machine-learning algorithm, the predictions may not only capture phenotypic characteristics beyond skin tone (e.g., facial morphology)

[10] About half of all candidate photos in our estimation sample are in grayscale. In Appendix C.2, we also assess the measurement quality across grayscale and colored photos, providing evidence that motivates the pooling of both types of photos in our main analysis.

Figure 2: Skin Tone Distribution

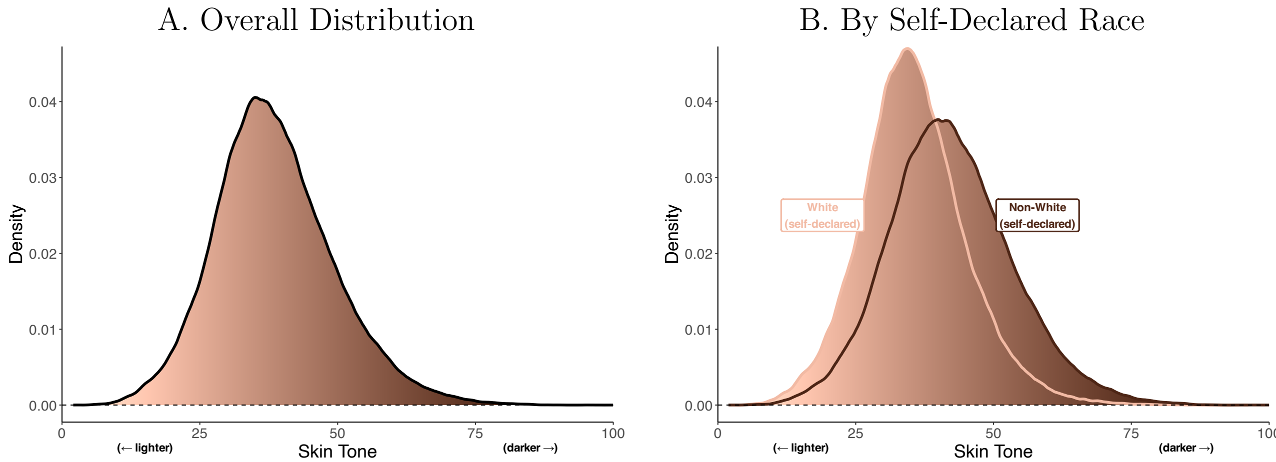

Notes: The figure shows the distribution of skin tone as measured from the electoral photos of candidates. As will be done in all figures in paper, the skin tone distribution is visualized by varying the L^* value of a skin tone with $a^* = 17.2$ and $b^* = 19.1$, which correspond to the sample averages of a^* and b^* across all our colored photos. Panel A depicts the skin tone distribution among all candidates. Panel B splits candidates by their self-declared race, distinguishing White and non-White (*pretos* and *pardos*) candidates.

but also more manipulable features that predict racialization by employers (e.g., grooming or wearing jewelry).

Estimation Sample. Our estimation sample consists of approximately 330,000 unique individuals who appear as political candidates at least once between 2014 and 2024 and who can be linked to at least one formal work contract in RAIS. For these matched individuals, we construct a yearly panel capturing all jobs of an individual between 2003 and 2015. We focus on private-sector, full-time ($\geq$ 30 hours/week), open-ended contracts held by individuals aged 20 to 54. For each year, we retain the primary job, defined as the contract with the highest annual earnings.

Our analysis compares employer-ascribed race in these jobs to the candidate's self-declared race and skin tone measured at the time of candidacy.[11] If an individual has run for a political office multiple times, we assign their modal self-declared race (across 2014-2024) and their mean skin tone (across 2004-2024). This approach implicitly treats self-identity and skin tone as time-invariant individual characteristics, which we compare to potentially divergent racial perceptions of different employers. In robustness checks, we relax this assumption and restrict comparisons to periods in which the different racial identity measures overlap in time.

Appendix Table B.1 presents summary statistics for our estimation sample of political candidates and compares them with the full RAIS workforce. Relative to other workers, candidates are more likely to be male (by 13 percentage points) and slightly older (by 0.84 years), but they are similar in education and average earnings. Moreover, as shown in Appendix Table B.2, the wage gap between White and non-White workers (based on employer-ascribed race) is 10 percentage points smaller among candidates than among other workers, suggest-

[11] We exclude observations where either the employer-ascribed or self-declared race is Asian or Indigenous, following standard practice in the literature (e.g., Gerard et al., 2021).

ing that our estimated disparities likely represent lower bounds for population-level gaps. Although political candidates are not representative of the Brazilian workforce, this sample is uniquely suited for our purposes because it allows us to compare three distinct measures of racial identity: race as perceived by the employer, self-declared racial identity, and skin tone as a phenotypic marker.

4 Mismatched Racial Identities

In this section, we document how employer-ascribed race differs from self-declared race and how the same worker is classified across employers. Throughout, we focus on the binary distinction between White and non-White. Appendix Table B.4 also reports results when disaggregating the non-White category into *pretos* and *pardos*.

Self-Declared versus Employer-Ascribed Race. Panel A of Table 1 cross-tabulates self-declared and employer-ascribed race. Roughly 72% of workers fall on the diagonal, that is, they self-identify and are classified by their employer as belonging to the same racial group (White or non-White). The remaining 28% exhibit a mismatch between self-identification and employer-ascription. The mismatch is asymmetric. Among self-declared Whites, about 18% are recorded as non-White by their employer. Among self-declared non-Whites, about 40% are reported as White. In other words, conditional on how workers describe themselves, employers are much more likely to "whiten" non-White workers than to "darken" self-identified Whites.

Skin tone helps to interpret these discrepancies. Within each cell of Panel A, we report the mean standardized skin tone measure. Workers who self-identify as non-White and are also classified as non-White have, on average, the darkest skin tones. Self-identified Whites who are classified as White have the lightest skin tones. The two off-diagonal cells, where self-identification and employer-ascription differ, are characterized by intermediate skin tone levels. Panel (a) of Appendix Figure A.3 shows the full skin tone distribution, confirming that mismatches concentrate among workers with mid-range skin tones rather than at the extremes of the distribution. Notably, the average tendency of employers to "whiten" rather than "darken" their workers does not appear to be driven by skin tone differences: self-identified non-Whites classified as White by their employer have skin tones that are 0.30 SD darker than those of self-identified Whites classified as non-White.

Taken together, these facts show that self- and employer-perceptions are tightly linked but not interchangeable. For many workers, especially those with ambiguous skin tone, the racial category encoded in administrative data is not a mechanical reflection of their self-understanding.

Mismatch across Employers. We next examine how the same individual is classified across employers. Panel B of Table 1 restricts the sample to workers who change jobs and compares employer-ascribed race in the previous and current establishment across consecutive years. About 77% of job switchers are assigned the same racial category by both employers. The remaining 23% change racial classification when they move. As in Panel A, the off-diagonal cells are sizable on both sides. Roughly 30% of workers who were previously

Table 1: Mismatch in Racial Identities

[A] Mismatch between Self-Declared and Employer-Ascribed Race

	Employer-Ascribed Race			
	White		Non-White	
Self-Declared Race	%	Avg. Skin Tone	%	Avg. Skin Tone
White	45.20%	-0.32	9.71%	-0.11
Non-White	17.84%	0.19	27.25%	0.45

[B] Mismatch across Subsequent Employers

	Employer-Ascribed Race, Current			
	White		Non-White	
Employer-Ascribed Race, Previous	%	Avg. Skin Tone	%	Avg. Skin Tone
White	49.82%	-0.22	12.18%	0.15
Non-White	10.88%	0.15	27.12%	0.40

Notes: The table displays the (mis)match shares for different race measures in our sample, as well as the average skin tone for each cell. Skin tone is standardized to have a mean of zero and a standard deviation of one across the whole sample (higher values represent darker shades, while lower values represent lighter shades). Panel A presents the distribution of workers by self-declared race (rows) and employer-ascribed race (columns). It considers all worker-year observations ($N = 1,510,833$). Panel B presents the distribution of ascribed race for the same worker across different employers in consecutive years. It considers all workers who changed jobs ($N = 153,366$) and presents the distribution of workers by their ascribed race in the previous employer (rows) and current employer (columns).

classified as non-White are recorded as White by their new employer, while close to 20% of previously White workers are coded as non-White in their subsequent job. Panel (b) of Appendix Figure A.3 shows that mismatches across employers are again concentrated among individuals with intermediate skin tones.[12]

These patterns indicate that employer-ascribed race is not a fixed attribute of the worker but the outcome of an interaction between the worker's appearance and the classification practices of each firm. To assess the role of firms' racial classification practices in more detail, we exploit workers' moves across firms in the full RAIS sample. Specifically, we run an AKM model (Abowd et al., 1999) that regresses an indicator for being ascribed non-White by the current employer on a full set of worker and firm fixed effects. This allows us to separate variation in firm-specific practices from variation in workers' portable racial cues. Appendix D provides details on the empirical specification, tests of the identification assumptions, and estimation results. The variance decomposition shows that only about half (47.2%) of the variation in employer-ascribed race is explained by worker effects – i.e., characteristics workers carry with them across jobs – while firm effects account for almost a third (29.6%) of the variation (see Appendix Table D.1).[13] The magnitude of

[12] While Table 1 focuses on the matched sample of political candidates, Appendix Table B.5 shows that the mismatch rates across employers are very similar within the full sample of workers in RAIS, suggesting that the candidate sample is not atypical in terms of racial ambiguity.

[13] The remaining variation is largely captured by the residual, given a very small covariance between worker and firm effects. When correcting for limited mobility bias in a leave-one-out connected set following the

firms' contribution is large, even higher than their role in explaining the variance of wages in the Brazilian labor market (Gerard et al., 2021). This underscores that firms play a quantitatively important role in the racial ascription of workers: some firms systematically "whiten" their workforce, whereas others tend to "darken" them.

Interpreting Mismatch and Measurement Concerns. A natural concern is that the observed mismatches may simply reflect noisy reporting rather than differences in latent race concepts – for instance, if employers pay little attention to reporting in RAIS. In the following, we discuss several pieces of evidence indicating that measurement error alone is unlikely to account for the observed mismatch.

First, as noted above, the mismatch is concentrated among intermediate skin tones, suggesting that self- and employer-ascribed racial categories diverge when phenotypic features allow for ambiguity. This is inconsistent with a scenario whereby employers record race with noise uniformly across the skin tone distribution. Instead, it suggests systematic divergence in ascribed race in cases where visual cues are ambiguous.

Second, we show that the mismatch patterns are similar across sub-samples that should exhibit higher reporting quality. Appendix Table B.5 repeats the comparison between self-declared and employer-ascribed race under a series of restrictions. We begin by excluding firms that uniformly assign the same race to their entire workforce.[14] Because these firms might plausibly be less attentive in their reporting, we re-estimate mismatch rates in the subset of establishments where we observe workers assigned to different racial groups. We then restrict the sample to firms with dedicated human-resources staff to address concerns that variation in classification firms' capacity to collect information on workers. We also limit the sample to worker-year observations in which both workers and employers report race in the same period, ruling out that discrepancies arise from changes in self-declared race over time. Finally, we focus on candidates for local councilor positions, among whom potential strategic incentives to report particular self-identified race are arguably weaker than for higher political offices. Across all these robustness checks, both the magnitude and the asymmetry of mismatches remain virtually unchanged. As an additional benchmark, the same table reports mismatches in gender, showing that disagreement between self-declared and employer-reported gender and changes in reported gender across employers are negligible. If employers filled out workers' demographic information in RAIS with large inattention, we would expect to observe similar mismatches in gender reporting, which we do not.

Third, we exploit the presence of multiple employer observations for a sub-sample of workers and aggregate employer-ascribed race across jobs. Appendix Table B.6 compares several measures of race perception – including the modal employer classification and the share of employers that code the worker as non-White – to self-declared race, following the standard approach in studies with RAIS data (e.g., Gerard et al., 2021). To address

approach of Kline et al. (2020), the share of variation explained by firm effects remains high (25.5%). The share explained by worker effects is reduced more substantially (to 30.8%,) as we estimate a higher correlation between firm and worker effects, accounting for 9.1% of the variation in ascribed race (see Column (2) of Appendix Table D.1).

[14] Most establishments do not assign the same racial category to all their workers: Appendix Figure A.1 shows that less than 8% of all workers are employed in fully-segregated establishments that classify their entire workforce as exclusively White.

measurement error in these aggregated measures, we estimate split-sample IV regressions in which the measure constructed from one random half of a worker's yearly observations is instrumented with the measure constructed from the other half. Even after averaging across employers and applying measurement error corrections, we find that the association between employer-ascribed and self-declared race remains far from one-to-one.

Overall, the results suggest that measurement error is unlikely to be the sole explanation for the large mismatch rates we document – 28% between self-declared and employer-ascribed race and 23% in ascribed race across subsequent employers. Most of the previous research has documented the fluidity of race by studying changes in self-identification over time (e.g., Muniz et al., 2024; Janusz, 2021; Liebler et al., 2017; Miranda, 2015; Saperstein and Penner, 2012). In our sample of political candidates, among those who ran for office multiple times between 2014 and 2024, self-declared race changes between White and non-White across subsequent elections in 16.9% of cases. Thus, the divergences between self- and hetero-identification, and across different observers, are substantially more common than changes in self-identification over time.

5 Race Function

The mismatch patterns documented in the previous section suggest that race can be ambiguous and that it is the outcome of a classification rule employers apply to observable worker characteristics. In this section, we characterize this employer's "race function." We first examine the probability that a worker is coded as non-White given continuous skin tone, both across the full sample and within self-declared racial groups. We then show how this mapping shifts with other worker characteristics, making explicit how employers construct racial categories from workers' skin tone, self-identity, socio-economic status, as well as the local context.

Skin Tone and Self-Identification. Figure 3 shows that skin tone is a central input into employers' racial ascription. Panel (a) shows a binned scatterplot of the probability that a worker is coded as non-White against skin tone. The relationship is steep and monotone: darker-skinned workers are much more likely to be classified as non-White, while very light-skinned workers are usually classified as White. Moving from the 10th to the 90th percentile of skin tone increases the likelihood of being ascribed non-White by 32 percentage points. Panel (b) splits the sample by self-declared race and shows that this pattern holds even within self-identification. Among workers who identify as White, the likelihood of being ascribed non-White rises with darker skin tone. Among workers who self-identify as non-White, lighter-skinned individuals are less likely to be coded as non-White. Employers, therefore, do not simply reproduce how workers describe themselves, but they incorporate their own visual impressions into racial classification.

At the same time, skin tone is far from a complete input of the race function. Both panels reveal substantial heterogeneity in the tails of the distribution. About 20% of the lightest-skinned workers are still classified as non-White, while over 30% of the darkest-skinned workers are coded as White. The patterns suggest that other worker characteristics

Figure 3: Relationship between Employer-Ascribed Race and Skin Tone

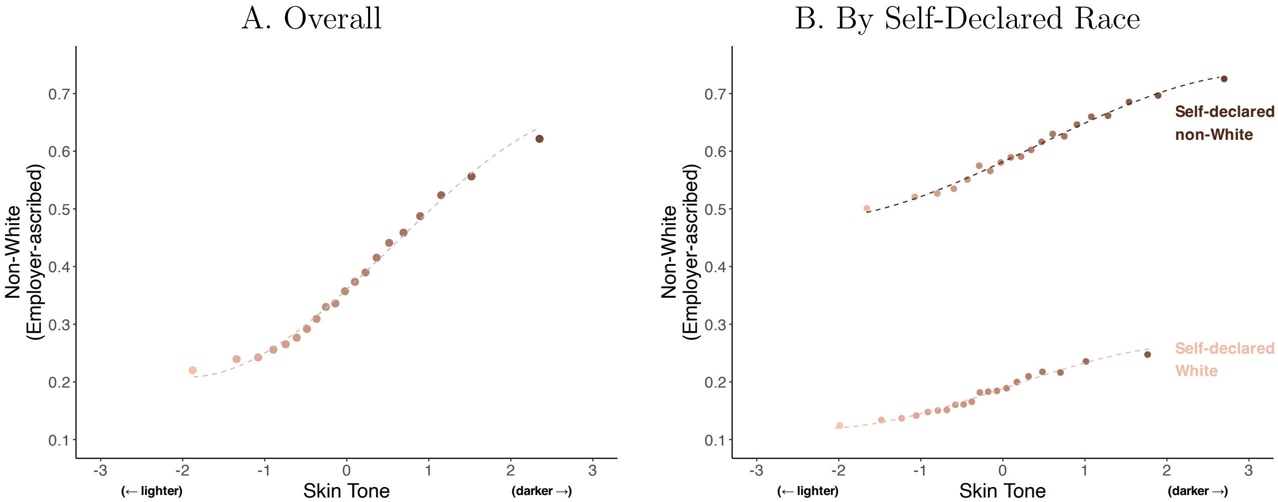

Notes: The figures show binned scatterplots of employer-ascribed race (an indicator equal to one if the worker is ascribed non-White) and workers' skin tone (standardized), both unconditionally (Panel A) and splitting the sample by self-declared race (Panel B). The plots are constructed following the procedure in Cattaneo et al. (2024). For each sample, we divide the skin tone distribution in twenty equal-sized bins and, for each bin, plot the average of employer-ascribed race. We fit third-degree polynomials based on the underlying data of each sample.

also shape racial ascription even at a given level of skin tone.[15]

Education and Context. Next, we illustrate that worker characteristics and local social context systematically shift how employers map skin tone into racial categories along two central dimensions: education and region. These two dimensions are particularly relevant because education is the most obvious (and measurable) productivity-relevant factor, and because Brazil is a country with immense racial diversity and significant heterogeneity in social norms around racial classification. Figure 4 examines how these factors reshape the race-skin tone relation, considering both its level and its slope. The level captures who is considered White or non-White at a given skin tone, while the slope captures how strongly employers translate incremental skin tone differences into racial categories. Together, these two components reveal not only who is coded as what, but also how sensitive racial ascription is to variations in phenotype.

Panel (a) shows how education changes the race function. Among workers with at most a high school education, the probability of being classified as non-White rises strongly with darker skin tone. For university-educated workers, this likelihood is lower across the whole skin tone distribution. More educated workers are therefore less likely to be coded as non-White than workers of the same skin tone who have less schooling. In addition, the slope between employer-ascribed race and skin tone appears smaller among workers with a university education. Thus, education not only raises the likelihood of being coded as White at any given level of skin tone (a lower threshold for Whiteness) but also attenuates the weight

[15]Alternatively, measurement error from our skin tone algorithm could explain some of the residual variation. We discuss measurement error in skin tone in more detail in Section 6.2.

Figure 4: Relationship between Employer-Ascribed Race and Skin Tone, by Education and Region

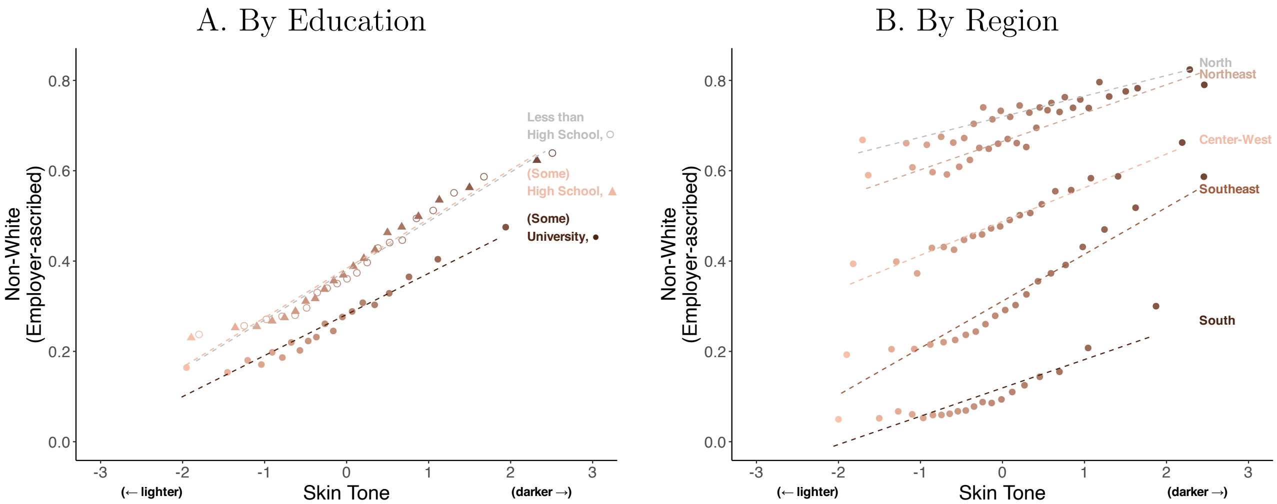

Notes: The figures show binned scatterplots of employer-ascribed race (an indicator equal to one if the worker is ascribed non-White) and workers' skin tone (standardized), splitting the sample by workers' educational attainment (Panel A) and region of the establishment (Panel B). The plots are constructed following the procedure in Cattaneo et al. (2024). For each sample, we divide the skin tone distribution in twenty equal-sized bins and, for each bin, plot the average of employer-ascribed race. We fit third-degree polynomials based on the underlying data of each sample.

employers place on skin tone (a flatter gradient).

Panel (b) shows that regional context also alters the mapping from skin tone to race. For a given skin tone, employers in the North classify workers as non-White at substantially higher rates than employers in the South, with the Northeast, Center-West, and Southeast in between. This suggests that "Whiteness" includes individuals with much darker skin tones in the South compared to the North. This classification practice reflects the local racial context: in the North, where the Afro-descendant population is largest, the racial norm is non-White and only very light-skinned workers stand out as White. In the South, shaped by European immigration, the racial norm is White, and only very dark-skinned workers are classified as non-White. The curves also differ in slopes: the Southeast exhibits a steeper gradient, suggesting that small tone differences carry more weight in racial ascription. Again, this is in line with the local racial context in the Southeast, characterized by a highly ethnically diverse population and thus sensitivity to racial categories. Together, both panels underscore that racial ascription in Brazil's labor market is not a mere reflection of phenotype but is influenced by worker characteristics and local racial context, which jointly shape employers' interpretations of the same underlying racial cue.

Multivariate Race Function. To quantify how employers translate multiple worker attributes into racial classifications, we estimate a multivariate "race function" using a linear probability model. Table 2 reports coefficients relating employer-ascribed race to a broad set of worker and labor market characteristics. All specifications include year fixed effects to absorb changes in racial coding across all worker-employer pairs over time. Our goal is twofold: to isolate the direct effects of these characteristics and to examine how the coefficient on skin tone evolves as richer controls and fixed effects are added.

We begin with the direct effects. Column 1 reports the unconditional association between skin tone and employer-ascribed race: a one-standard-deviation increase in skin tone raises the probability of being classified as non-White by 11.0 percentage points. Notably, skin tone explains less than 6% of the variation in employer ascriptions. highlighting the strong potential for other characteristics to determine how employers perceive the racial identity of their workers. Column 2 introduces demographic characteristics, that is, education, gender, and a squared age term. Conditional on skin tone, university graduates are 9.6 pp less likely to be coded as non-White than those with less than a high school education, which confirms that status signals shape race perceptions. Column 3 adds regional fixed effects. Relative to the North and the Northeast, where the shares of Afro-descendant and Indigenous populations are highest, employers in all other regions whiten workers more readily at a given skin tone.

Columns 4 to 6 introduce progressively richer labor-market fixed effects, which allow us to compare workers who face similar task requirements, pay scales, and employer environments but differ in phenotypic and demographic attributes. Column 4 includes industry fixed effects; column 5 adds three-digit occupation by industry fixed effects.[16] Column 6 includes establishment-by-three-digit-occupation fixed effects, effectively holding the employer, i.e., the interpreter of race, constant and exploiting variation in phenotypic and status cues among workers that hold the same job. Across all these specifications, the qualitative pattern remains: education decreases the likelihood of being classified as non-White, conditional on skin tone, even within narrowly defined labor-market segments.

These results support a central claim of the constructivist approach: productivity-relevant characteristics shape race perceptions. To substantiate this claim further, we leverage information from previous employers for a sub-sample of workers who hold multiple jobs over our observation period. Appendix Table B.7 replicates Table 2 but adds labor market seniority (tenure conditional on age) and the employee's final wage at their previous job to the race function. Seniority reduces the likelihood of being classified as non-White, and higher wages at the previous firm are also associated with whitening, providing additional evidence of the productivity-racialization path.

Turning again to the effect of skin tone in Table 2, we find that darker-skinned individuals remain more likely to be classified as non-White across all specifications, but the magnitude of the skin tone coefficient decreases by approximately 40% in column 6 compared to column 1. The skin tone elasticity attenuates as individuals with varying skin tones differ in worker and job characteristics that also matter for employers' race perceptions. Thus, once we compare a dark-skinned and a light-skinned individual with the same characteristics, the perceived identity contrast between them is far less pronounced than the unconditional contrast between two individuals drawn at random from the population. Moreover, the results from column 6 indicate that workers with varying skin tones are sorted differently into more or less whitening employers, which explains the attenuation of the skin tone coefficient relative to column 5 (by 13%). At the same time, the results also show that skin tone remains a significant predictor even within job cells, highlighting that racial categorization is specific

[16] We consider 9 industry categories following the Brazilian national classification of economic activities: agriculture, extractive industries, manufacturing, construction, retail, transport & communication, finance & real estate, education & health, and other services. Three-digit occupations correspond to the 192 sub-categories in the Brazilian Occupation Classification (2002 CBO).

Table 2: The Race Function: Determinants of Employer-Ascribed Race

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Skin Tone	0.110***	0.107***	0.082***	0.081***	0.078***	0.068***	0.041***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							0.235***
							(0.003)
Demographics							
Male		0.002	-0.004**	-0.011***	-0.013***	-0.014***	-0.012***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Age		-0.025***	-0.014***	-0.016***	-0.013***	-0.012***	-0.009***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age^2		-0.007***	-0.003***	-0.003***	-0.002***	-0.002**	-0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
(Some) High School		-0.009***	-0.034***	-0.027***	-0.012***	-0.007**	-0.001
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
(Some) University		-0.096***	-0.115***	-0.110***	-0.058***	-0.032***	-0.014**
		(0.003)	(0.003)	(0.003)	(0.003)	(0.006)	(0.006)
Region							
Northeast			-0.052***	-0.053***	-0.055***		
			(0.004)	(0.004)	(0.004)		
Southeast			-0.400***	-0.398***	-0.399***		
			(0.003)	(0.003)	(0.003)		
South			-0.584***	-0.578***	-0.578***		
			(0.003)	(0.003)	(0.003)		
Center-West			-0.226***	-0.224***	-0.223***		
			(0.004)	(0.004)	(0.004)		
Constant	0.370***	0.392***	0.752***	0.750***	0.737***	0.390***	0.277***
	(0.001)	(0.002)	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)
N	1,510,833	1,510,833	1,510,833	1,510,510	1,510,445	1,249,942	1,249,942
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes
R-squared	0.057	0.065	0.209	0.213	0.223	0.779	0.793
Adjusted R-squared	0.057	0.065	0.209	0.213	0.222	0.719	0.736

Notes: The table reports OLS estimates of the relationship between employer-ascribed race (an indicator equal to one for non-White) and various covariates. The first column displays the relationship with skin tone, when only controlling for year fixed effects. The subsequent columns progressively add demographic covariates as well as region, one-digit industry, three-digit occupation, and firm fixed effects. The last column also conditions on individuals' self-declared race. Standard errors clustered at the individual level are in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

to worker-employer pairs.

Finally, column 7 introduces self-declared race. Although self-identification may itself be influenced by the same factors that shape others' perceptions, and may therefore be a "bad control," our goal here is to hold constant workers' own identity and examine the remaining role of phenotype and demographics in generating divergence between employer and worker perceptions.[17] As expected, self-identifying as non-White increases the likelihood

[17] We estimate a separate race function for self-identification in Appendix Tables B.8 and B.9. The results indicate that self-identification places different weights on observable characteristics, further emphasizing that self-identification and employer perception are related but distinct outputs of different race functions.

of being coded as such by employers, either because workers and employers interpret the same underlying cues similarly or because workers signal their identity in ways that employers (partially) incorporate. The coefficient on skin tone falls by about 40% once self-identification is included, consistent with shared underlying determinants of racial identity. Yet, even when holding self-identification constant, employer-ascribed race is affected by skin tone and by other social cues, such as education. This reinforces the conclusion that self-identification and employer perception are not the same; in fact, employer-ascribed race still responds to phenotypic and social cues, even conditional on self-identity.

In summary, Table 2 shows that employers' racial perceptions are related to various worker characteristics and local context, among workers with the same skin tone. For example, workers with a university education are more likely to be perceived as White. We emphasize that this association need not be interpreted as capturing the causal effect of an exogenous increase in education. Rather, education may proxy for a broader bundle of socioeconomic characteristics that vary with education, and that may shape how employers racialize workers, such as income, dress, or speech patterns. Under this interpretation, the substantive implication is unchanged: racial ascription responds to social-status cues, going beyond purely phenotypical characteristics.

Machine-Learning Prediction of Visual Cues. Skin tone is not the only phenotypic trait that employers may use as a racial marker. Facial morphology, hair texture, eye color, and other visible characteristics may affect racial ascription and may be only imperfectly captured by our skin tone measure. In addition, some of these traits may also be correlated with socioeconomic characteristics, raising the concern that the education gradient documented above partly reflects omitted phenotypic cues rather than social information beyond appearance.

To examine these issues, we account for the photo-based machine-learning predictions of employers' racial ascriptions introduced in Section 3.2 and described in more detail in Appendix C.3. These predictions provide a flexible summary of the cues visible in candidates' photos that are associated with employer ascription. Importantly, this black-box machine-learning model may not only capture phenotypic traits beyond the skin-tone measure, but also more manipulable features, such as hairstyle, grooming, visible oral health, jewelry, and photo quality. Because these features may themselves reflect socioeconomic status, conditioning on the machine-learning prediction is a demanding test of whether social-status cues matter: it absorbs not only phenotypic traits, but also any component of socioeconomic presentation that is visible in the photo.

Appendix Table B.10 shows that the main insights from Table 2 remain similar when we replace skin tone with the machine-learning predicted probability of employer non-White ascription. The predictions explain about 17% of the variation in employers' racial ascriptions, roughly three times the explanatory power of skin tone alone. Crucially, education remains significantly associated with ascribed race even after controlling for the flexible machine-learning prediction from photos, although the coefficients are attenuated. For example, in column 2, workers with university education are 4.4 percentage points less likely to be perceived as non-White, compared to 9.6 percentage points when conditioning only on skin tone.[18]

Overall, this test reinforces the conclusion that racial ascription is shaped by social-status cues. It mitigates the concern that the patterns in Table 2 are driven only by visual traits beyond skin tone. For the remainder of the paper, we focus on the algorithmic measure of skin tone, which is more difficult to manipulate than other visual cues captured in photographs. It also provides a transparent and interpretable metric of an important phenotype dimension in Brazil (Telles, 2014) and allows our results to be readily compared with studies of colorism.

6 Racial Wage Disparities

Our results so far are consistent with a constructivist view of race: racial identities diverge between workers and employers and across employers, and racial ascription is a function of socio-economic cues. In the following, we examine the implications for measured disparities. In a first step, we present three estimates of racial wage disparities using employer-ascribed race, skin tone, and a perception-normalization approach. In a second step, we highlight the importance of modeling race perceptions by examining *i*) conditional versus unconditional contrasts *ii*) heterogeneity across workers and *iii*) the role of measurement error. Lastly, we discuss the interpretation and limitations of our results.

6.1 Estimates of Racial Disparity

Conventional Disparity. We start by quantifying racial wage gaps that are commonly estimated using RAIS data. Panel A of Table 3 shows differences in log hourly wages between workers who are ascribed White vs. non-White by their employer.[19] In column 1, conditioning only on year fixed effects, workers classified as non-White earn 14.8% less than workers classified as White.[20] This unconditional gap captures the disparate impact of race which entails both direct and systemic forms of discrimination (Bohren et al., 2025b).

Researchers often run "benchmarking regressions" to account for productivity differences among workers. They seek to quantify direct discrimination (both taste-based and statistical), understood as an employer's disparate treatment based on race. Columns 2 to 6 implement a standard benchmarking sequence, adding the same covariates and fixed effects used in our race-function analysis (see Table 2). The coefficient attenuates substantially when we control for sociodemographics and region fixed effects in columns 2 and 3, and shrinks further once we hold fixed sorting across industries and occupations in columns 4 and 5. The most stringent specification in column 6 compares workers within the same firm

[18] Only in column 7, where we control for self-declared race, education loses significance, potentially because other (more manipulable) traits correlate with self-declared race and education. In contrast, Appendix Table B.11 shows that labor market seniority and previous earnings whiten workers even conditional on self-declared race, the full set of visual cues and within firm-by-three-digit occupation cells.

[19] Appendix Table B.12 and B.13 restrict the sample to individuals running for city councilor positions (thus excluding high-profile offices) and to years before the candidate ever ran for office. These restrictions ensure strong labor market attachment of the sample and potentially lower incentives for public strategic signaling of race.

[20] Appendix Table B.14 expresses wage gaps in terms of workers' self-declared race. The unconditional wage gap for self-declared race lies at 16.3% and is thus modestly larger than for employer-ascribed race.

Table 3: Wage Disparities

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
[A] OLS: Conventional Disparity							
Non-White Employer-Ascribed	-0.148***	-0.103***	-0.054***	-0.064***	-0.038***	-0.021***	-0.017***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Non-White Self-Declared							-0.015***
							(0.002)
[B] Reduced Form: Colorism Disparity							
Skin Tone	-0.026***	-0.032***	-0.026***	-0.028***	-0.015***	-0.005***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							-0.017***
							(0.002)
[C] First Stage: Race Function							
Skin Tone	0.110***	0.107***	0.082***	0.081***	0.078***	0.068***	0.041***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							0.235***
							(0.003)
[D] 2SLS: Perception-Normalized Disparity							
Non-White Employer-Ascribed	-0.235***	-0.301***	-0.314***	-0.348***	-0.193***	-0.080***	-0.084***
	(0.011)	(0.010)	(0.014)	(0.014)	(0.012)	(0.013)	(0.022)
Non-White Self-Declared							0.003
							(0.006)
N	1,510,833	1,510,833	1,510,833	1,510,510	1,510,445	1,249,942	1,249,942
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	Yes	Yes	Yes	No	No
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes

Notes: The table reports different estimates of racial wage gaps based on employer-ascribed race. All regressions include year fixed-effects. The first column always shows the unconditional relationship (apart from year fixed-effects), and the remaining columns progressively add controls. The final column includes all fixed-effects, as well as controls for self-declared race. Panel A shows coefficient estimates of OLS regressions of log hourly wage on employer-ascribed race. Panel B shows reduced-form regressions of log hourly wage on skin tone. Panel C presents the first stage coefficient from regressing employer-ascribed race on skin tone. Finally, Panel D shows 2SLS estimates of regressing log hourly wage on employer-ascribed race, instrumented by skin tone. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

and three-digit occupation; even within these narrow cells, being classified as non-White is associated with a 2.1% lower wage.

Column 7 then distinguishes perception from self-identity by including both employer-ascribed and self-declared race. Both matter: self-declaring non-White is associated with 1.5% lower wages, while being classified as non-White by the employer carries an additional 1.7% penalty. This means that two workers who both self-identify as White and work in the same firm and occupation can earn different wages if one is perceived as non-White by the employer, directly connecting the mismatch documented in Section 4 to wage-setting.

The large attenuation across columns is not, by itself, evidence that "discrimination disappears" once we benchmark. Rather, controls also partial out determinants of racial classification itself. Therefore, it may not be obvious what the residual variation in race perceptions entails and how this changes the interpretation of the conventional disparity estimates. We return to this issue in Section 6.2.

Colorism Disparity. An alternative approach that is more explicit about the exploited variation focuses on one specific racial cue that enters employers' race function. We rely on our algorithmic measure of skin tone which shifts the likelihood that an employer classifies a worker as non-White even within narrowly defined cells, as shown in Section 5.

Figure 5 displays a binscatter plot of skin tone and log hourly wages, revealing a near-monotone negative relationship between the two. Panel B of Table 3 quantifies this gradient in a linear regression. In column 1, a one standard deviation darker skin tone is associated with 2.6% lower wages. This estimate is essentially unchanged when adding demographics, region fixed effects, and industry fixed effects, but shrinks substantially when adding occupation by industry or occupation by firm fixed effects. In the most stringent specifications in columns 6 and 7, a one standard deviation darker skin tone is associated with 0.5% lower wages and the disparity persists even conditional on workers' self-declared race. These results are consistent with the findings of the colorism literature that has documented differences across labor market outcomes for individuals with varying skin tones (Woo-Mora, 2026; Abramitzky et al., 2023; Goldsmith et al., 2006) and shows that this continues to hold within firm-by-occupation cells.

While both conventional and colorism disparity estimates attenuate substantially across columns, their sensitivity to specific controls differs. The conventional wage gap in Panel A is most sensitive to regional controls, while the colorism disparity in Panel B is most sensitive to occupation and firm-by-occupation fixed effects, suggesting that sorting of workers operates differently along employer-perception and skin tone margins. Notably, drops in the reduced-form estimates in Panel B coincide with a reduced elasticity of employer-ascribed race with respect to skin tone in Panel C, suggesting that part of the decline in the colorism disparity is driven by a weaker racialization margin.[21]

Perception-Normalized Disparity. The colorism disparity provides an intent-to-treat effect that does not take into account how skin tone enters the race function. To account for how skin tone translates into race perceptions, in Panel D of Table 3, we estimate a 2SLS model that instruments employer-ascribed race with skin tone, thereby rescaling the colorism disparity by the corresponding race function.

Column 1 reports a large unconditional perception-normalized disparity of 23.5%. Adding demographics in column 2 and then region fixed effects in column 3 leaves the PND large and, if anything, increases it in magnitude. The estimate becomes largest in column 4 when we additionally include industry fixed effects, consistent with the idea that once we compare

[21] For instance, an about 20 percent drop in the reduced-form estimate occurs when adding region fixed effects, accompanied by a similar drop in magnitude in the race function. This is in line with our previous findings in Section 5, showing substantive heterogeneity in the racialization margin of skin tone across regions.

Figure 5: Relationship between Wages and Skin Tone

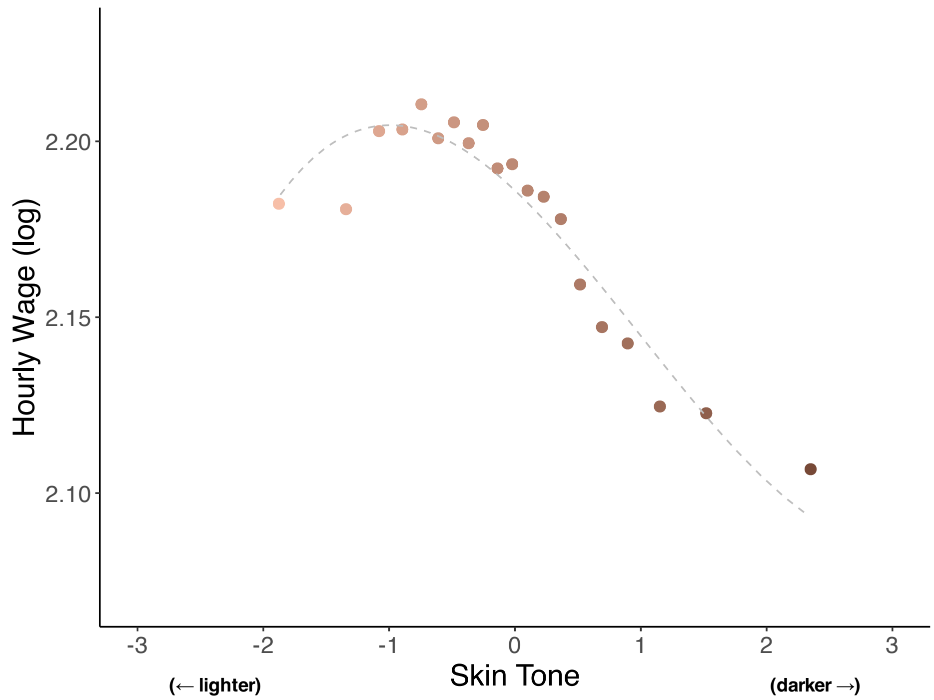

Notes: The figure shows a binned scatterplot of log hourly wage on worker's skin tone (standardized). The plot is constructed following the procedure in Cattaneo et al. (2024). We divide the skin tone distribution in twenty equal-sized bins and, for each bin, plot the average log hourly wage. We fit a third-degree polynomial based on the underlying data.

workers within broad sectors the remaining wage differences associated with skin-tone based perception are even larger. Introducing industry-by-occupation fixed effects in column 5 reduces the estimate considerably, indicating that a meaningful share of the wage gap operates through sorting into occupations. In the most stringent specification, which compares workers with the same demographics who work in the same job (column 6), we estimate a PND of 8.0%. Moreover, the results remain essentially unchanged when additionally controlling for workers' self-declared race in column 7, indicating that within job cells the wage penalty loads on how the employer classifies the worker rather than on self-identification. This result also rules out the PND estimates pick up unobserved factors that determine both self-declared race and wages.

6.2 The Importance of Measuring Race Perceptions

Conditional vs. Unconditional Contrasts. Racial disparity estimates based on employer-ascribed race and skin tone decline substantially as we add controls, whereas the PND falls much less. Panel A of Table 3 shows that the employer-ascribed race wage gap shrinks by 86% in the within-firm specification relative to the unconditional gap. Part of this decline reflects differences in productivity-related characteristics. But, as discussed in Section 2, part of it reflects netting out the racial identity contrast itself: the same covariates used to proxy for productivity also enter the racial perception process. Interpreting employer-ascribed race disparities therefore requires making explicit that the estimand depends on the residual variation in recorded race after conditioning on controls.

Panel B of Table 3 shows an analogous pattern for skin tone. The reduced-form colorism estimates shrink by about 81% when moving from column 1 to column 6. Under the constructivist framework, this decline can reflect two mechanisms simultaneously. Conditioning can make workers more comparable in productivity-relevant dimensions, but it can also make skin tone less predictive of employer-ascribed race. As shown in Table 2 (and replicated in Panel C of Table 3), the relationship between skin tone and employer-ascribed race falls by about 40% when moving from the unconditional specification to one with demographics and detailed firm-by-occupation fixed effects. The elasticity declines as skin tone is correlated with other worker and job characteristics that enter the employer's classification rule.

The PND approach explicitly takes into account the attenuation of the first stage. After scaling the reduced form by the corresponding change in perceptions, the PND declines by about 66% between columns 1 and 6 in Panel D – a decline that is about 20% smaller than that for the employer-ascribed or colorism disparity. These results highlight how conventional conditional disparity estimates can overstate the role of productivity differences (and other differences in worker and job characteristics) in explaining observed racial wage gaps. Moreover, they emphasize that measuring and incorporating the first stage of race perceptions is crucial to allow for differences in the racialization margin under conditional contrasts.

Heterogeneity across Worker Groups. Heterogeneity in cue-based disparity is not automatically evidence of heterogeneous discrimination. This is because the the same racial cue need not enter the employer's classification rule with the same weight for all workers. For instance, differences in reduced-form estimates by education can reflect either differences in the weight employers place on skin tone when assessing workers' race or actual differences in discrimination across education groups. Figures 6 and 7 illustrate this point for heterogeneity by education and region, respectively.[22] Each figure reports group-specific estimates of wage disparities using the relevant race measure, along with relative differences across groups (in percentages) and stars that indicate whether differences are statistically significant. The results come from a regression that includes interactions of group indicators with the relevant race measure, as well as interactions with all other controls and fixed effects in the respective specification.

Consider first heterogeneity by education in Figure 6. Panel B shows that reduced-form estimates of skin tone wage disparities tend to be larger for individuals with university education than for those with less than a high school degree, consistent with previous findings of larger racial wage gaps among highly-educated workers in Brazil (IBGE, 2022; Gerard et al., 2021). When controlling at least for basic demographics, the point estimates are 23-64% larger for the highly educated.[23] However, from our discussion of Figure 4 in Section 5, we

[22] Appendix Figures A.4 and A.5 report analogous exercises by gender and age, where we find little evidence of heterogeneity.

[23] In the first specification, which includes only year fixed effects, we find small and largely insignificant reduced-form estimates for all three education groups. The results change once we add controls for gender (and age) in the second specification, as darker skin tone is positively correlated with being male in our sample of political candidates (correlation coefficient = 0.197).

Figure 6: Wage Disparities, Heterogeneity by Education

A. OLS: Conventional Disparity

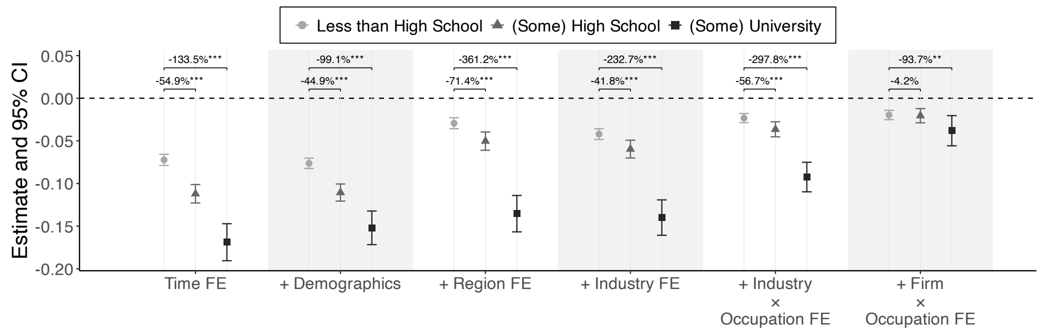

B. RF: Colorism Disparity

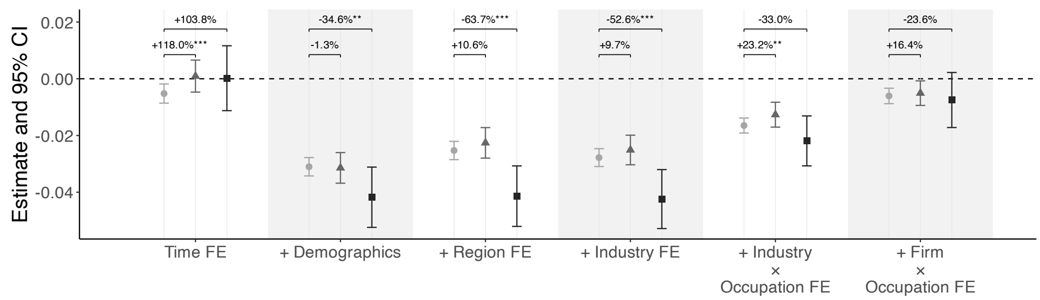

C. FS: Race Function

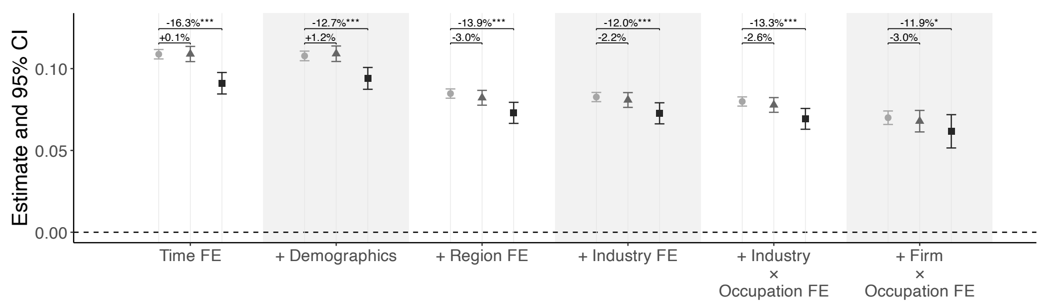

D. 2SLS: Perception-Normalized Disparity

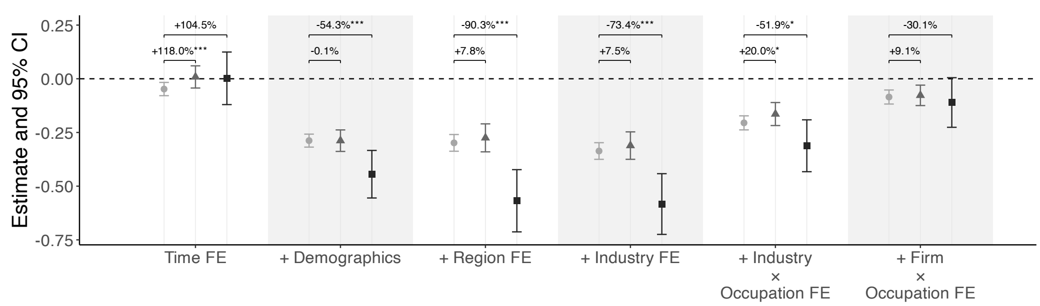

Notes: The figures display different estimates of racial wage gaps for different educational levels. We use the same specifications as columns (1) to (6) of Table 3, with interaction terms for the three educational levels. The plots present point estimates and 95% confidence intervals computed using clustered standard errors at the individual level. The brackets above pairs of point estimates report the percent difference between the estimates (relative to the absolute value of the reference category), with stars indicating whether the difference is statistically significant. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Figure 7: Wage Disparities, Heterogeneity by Region

A. OLS: Conventional Disparity

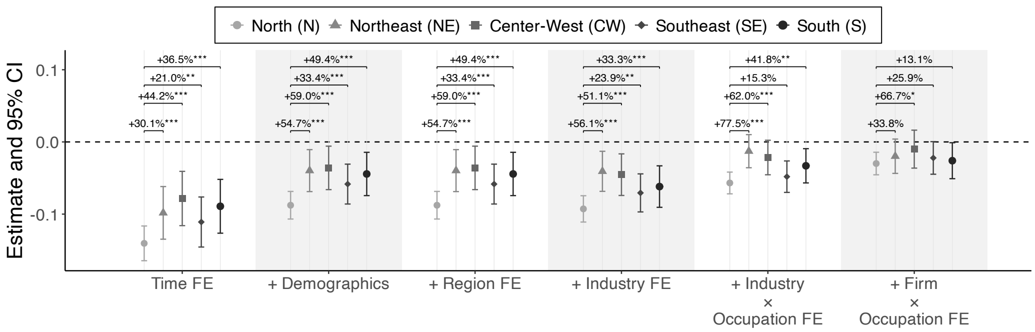

B. RF: Colorism Disparity

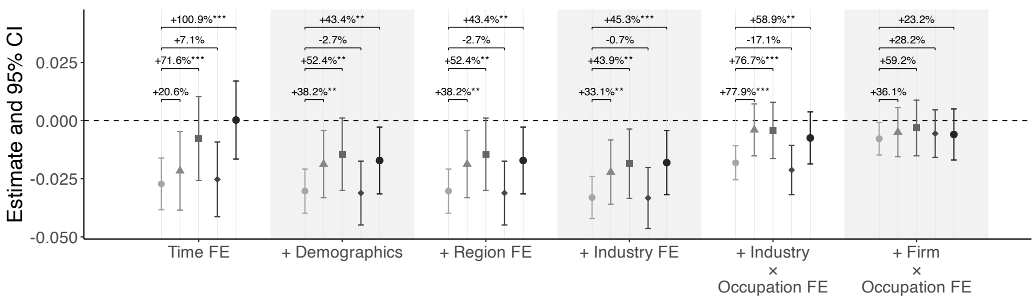

C. FS: Race Function

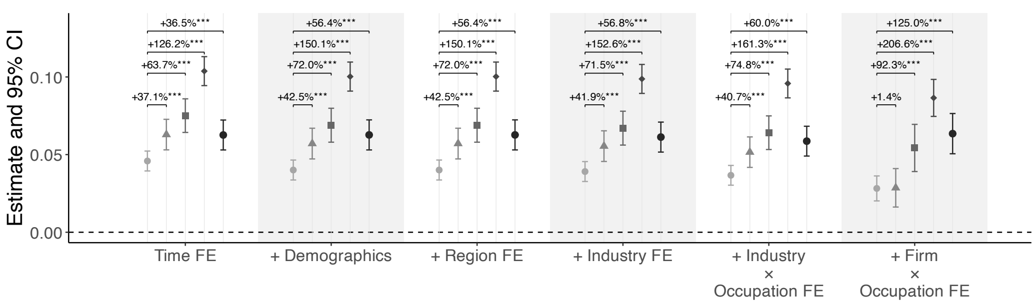

D. 2SLS: Perception-Normalized Disparity

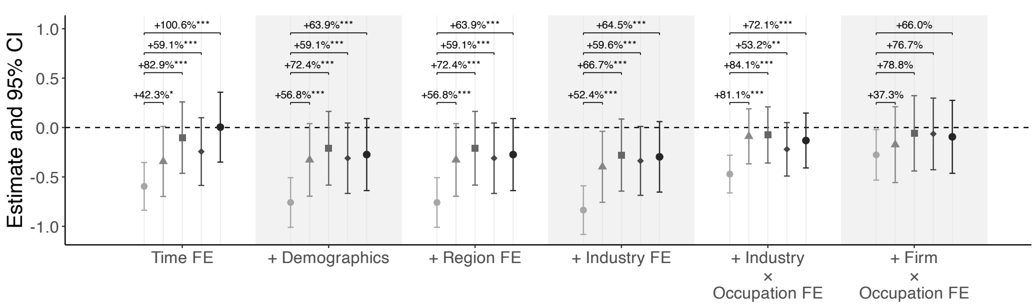

Notes: The figures display different estimates of racial wage gaps for the five different regions of the country. We use the same specifications as columns (1) to (6) of Table 3, with interaction terms for each region. The plots report point estimates and 95% confidence intervals, computed using clustered standard errors at the individual level. The brackets above pairs of point estimates report the percent difference between the estimates (relative to the absolute value of the reference category), with stars indicating whether the difference is statistically significant. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

know that the elasticity of employers' race perceptions with respect to skin tone is smaller among university-educated workers than among low-educated workers. Panel C of Figure 6 demonstrates that this fact holds even within firm-occupation cells. Throughout, skin tone is less predictive of employer-ascribed race for the highly educated compared to those with less schooling. Therefore, when considering the PND estimates shown in Panel D, which take into account differences in the skin tone elasticity, we get even larger gaps among university-educated workers than among low-educated ones. Specifically, among all specifications that control for basic demographics, the estimated PND for university-educated workers is 30-90% larger than for low-educated workers.

Regional heterogeneity, shown in Figure 7, reinforces this point. As an illustrative example, consider differences in wage disparities between the Southeast and the North. Reduced-form estimates based on skin tone alone indicate no significant differences in wage gaps across the two regions. However, the skin tone elasticity of employer-ascribed race is more than twice as high in the Southeast than in the North, a pattern that gets even stronger within stringent fixed-effect cells. Taking into account these differences in racialization, the PND estimates show that perception-normalized disparities are in fact significantly smaller in the Southeast than in the North.

Taken together, these results underscore that heterogeneity in reduced-form disparities need not reflect heterogeneity in discrimination. Differences in how strongly skin tone maps into perceived race can mask or exaggerate group-level comparisons unless the first stage is explicitly accounted for.

Measurement Error. Perhaps most obviously, racial wage gaps may be attenuated when race variables or racial cues are measured with noise. We quantify the role of measurement error in skin tone in attenuating colorism disparity, and then discuss how the PND approach helps with measurement error in both skin tone and employer-ascribed race.

We implement a split-sample IV design that uses individuals with multiple photographs from a subsample who ran for office at least twice. For each individual, we randomly split their photos into two halves, compute average skin tone in each half, and instrument the average from one half with the average from the other. Appendix Table B.15 reports results for the reduced-form effects of skin tone on wages and the first-stage effects on employers' race perceptions. We find a substantial role of measurement error in attenuating skin tone differences. The unconditional split-sample IV results indicate that a one-standard-deviation darker skin tone reduces wages by 11.4% (versus 5.1% in OLS) and increases the probability of being reported non-White by the employer by 34.6 percentage points (versus 15.2 percentage points in OLS). We find the same pattern across all specifications with controls: the split-sample IV coefficients are always more than twice as high as the corresponding OLS coefficients. Attenuation bias is therefore a serious concern for skin tone disparity estimates.

A useful feature of the perception-normalized disparity is its invariance to classical measurement error in skin tone. Noise in the skin tone measure attenuates the reduced-form and first-stage by the same factor, thus leaving their ratio (i.e., the PND) unchanged. Consistent with this, the PND implied by the OLS and split-sample IV estimates in Appendix Table B.15 is nearly identical (e.g., $-0.051/0.152 \approx -0.114/0.346$), despite large differences in the reduced-form and first-stage levels.[24]

The PND approach may also help bound the bias from measurement error in employer-ascribed race. On the one hand, measurement error may attenuate conventional OLS estimates in which ascribed race is the binary explanatory variable. On the other hand, PND estimates of disparities may be inflated because they scale reduced-form effects of skin tone by the first-stage effect of skin tone on employer ascription, and this first-stage relationship is itself attenuated by noise in the binary outcome (see, e.g., Bingley and Martinello (2017) on bias in IV estimators with a mis-measured binary endogenous regressor). Thus, if the gap between OLS and PND estimates were driven solely by measurement error in ascribed race, the true perception-based disparity is expected to lie between the two.

Overall, if the data-generating process gives reason to believe that measurement error in skin tone or race variables is a concern, then the PND approach can help recover an interpretation of disparities that avoids attenuation bias.

6.3 Interpretation and Limitations

Our results show that perception-normalized measures can differ substantially from conventional disparity estimates. In our most stringent specification, the PND is about four times as large as the wage gap computed directly from employer-ascribed race. This difference could be attributed to several factors: the endogeneity of employer perceptions, measurement error in perceptions, and the fact that the PND identifies a LATE. Therefore, in other settings, the OLS could exceed the PND, depending on the relative importance these forces. In fact, in a world where status and money "whiten" – and where none of the other forces are at play – racial disparity estimates based on endogenous perceptions would exceed the PND. The key point of our results is therefore not the magnitude itself but that the choice of racial measure and estimand is consequential and must be made explicit in a setting where race is not fixed but perceived.

The PND provides a scaling metric that incorporates the relevance of a racial cue in determining racial perceptions, thereby clarifying the interpretation of reduced form estimates. Heterogeneity in colorism estimates can arise not only from differences in perception-based discrimination but also from differences in how strongly skin tone maps onto racial perceptions. Scaling disparities in terms of perceptions allows for more meaningful comparisons of disparities across different groups, specifications, and measures of racial cues.

Importantly, interpreting the PND as the causal effect of employer-ascribed race on wages requires that the racial cue affects wages only through perceived race, and that it is as good as random conditional on controls. Even within narrowly defined cells, differences in skin tone may still correlate with productivity. Moreover, while observing employers' perceptions of workers' racial identity is a key strength of our setting, we lack information on perceptions held by other potentially relevant actors, in particular coworkers and customers. We interpret employer-recorded race in RAIS as a salient institutional perception, but it may also proxy correlated perceptions held by these other actors. In that case, the PND captures an employer-anchored, perception-normalized disparity rather than the causal effect of employer perception alone.

[24] They are not exactly equal because of finite-sample behavior of the SSIV.

More fundamentally, the exclusion restriction also requires that employers set wages based on perceived racial categories, rather than having direct preferences over workers with different skin tones.[25] A long tradition of research in social psychology and economics shows that individuals interpret, reason, and act through the lens of social categories (e.g., Charness and Chen, 2020; Shayo, 2020; Tajfel, 1981; Tajfel et al., 1971; Taylor et al., 1978). If, instead, disparities arise as a consequence of continuous skin tone differences, it remains unclear whether skin tone should be treated as a fixed worker trait, or whether employers' assessments of skin tone are themselves socially constructed in a manner analogous to perceived racial categories.[26]

Even if the exclusion restriction were satisfied, the resulting IV estimand would be local: the PND captures effects for workers whose ascribed race is shifted by the cue. This need not generalize across workers or other cues. We have documented that the first stage can differ across workers' education or local context.[27] A PND based on skin tone can differ from one based on names, hair texture, or dialect. In the Brazilian context, skin tone is considered the primary input into racial classification – more so than names, facial features, or other phenotypic cues (Telles, 2014). More generally, any racial cue must generate a sufficiently strong first stage to yield a meaningful PND.

7 Conclusion

This paper adopts a constructivist perspective on race in the labor market and shows that racial classification is not a fixed worker attribute but rather an outcome of perception shaped by phenotypic traits, socio-economic status cues, and context. A central lesson of our analysis is that racial classification is inherently perception-dependent. Different empirical approaches, therefore, target different objects, yet this distinction is rarely made explicit.

Racial disparity estimates require careful interpretation for at least five reasons. First, self-perceptions and decision-makers' perceptions can diverge. Second, recorded race may embed productivity-relevant signals. Third, benchmarking controls can absorb substantial variation in recorded race. Fourth, reduced-form estimates confound disparity with first-stage strength. Fifth, measurement error attenuates estimated gaps. Importantly, when the racial perception process is ignored, disparities are easily attributed to productivity differences that, in fact, partly reflect how workers are racially classified. Several limitations qualify the interpretation of our PND estimates, which should be viewed primarily as a

[25] Colorism studies document skin tone disparities among individuals within the same racial categories (e.g., Woo-Mora, 2026; Monk, 2016, 2014; Kreisman and Rangel, 2015; Goldsmith et al., 2006). Note, however, that if racial categories are endogenous objects, such conditional analyses also yield biased estimates of skin tone disparities (Adukia et al., 2025).

[26] In Appendix C.2, we analyze data on human raters' assessment of skin tone for a sub-sample of Brazilian political candidates collected by Bueno and Dunning (2017). We document substantial variation in how humans evaluate the skin color of the same candidate.

[27] To investigate the characteristics of the effective population in our PND regressions, Appendix Table B.16 shows estimates of the average characteristics of the effective sample following Hull (2025). We find that compliers (i.e., individuals for whom skin tone shifts employer-ascribed race) are overall similar to the full sample, though they are less likely to have university degrees and more likely to be located in the Southeast. This matches the race function results, where we found the elasticity of employer-ascribed race to skin tone to be larger among workers from the Southeast and among those with less educational attainment.

way to scale cue-based wage differentials by racial perception rather than as a clean causal estimate of discrimination.

This first large-scale empirical test of the constructivist view of race also points to several directions for future research. Our analysis focuses on measurement and interpretation in observational labor-market data, but similar issues arise in experimental designs. Audit studies or laboratory experiments could explicitly model and manipulate the cues that shape racial perception to study how decision-makers map observable characteristics into social categories and make decisions based on those categories. A second direction is to examine constructivist racial classification in other institutional settings where decisions rely on socially perceived categories, such as judicial sentencing, school admissions, or credit markets. The insights also extend to other social categories, including religion, caste or migrant status. Progress in the empirical study of inequality requires treating race, caste, or religion not only as a characteristic of individuals but also as a socially produced variable whose measurement depends on the underlying perception processes.

References

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Abramitzky, Ran, Jacob Conway, Roy Mill, and Luke Stein, "The Gendered Impacts of Perceived Skin Tone: Evidence from African-American Siblings in 1870–1940," *NBER Working Paper 31016*, 2023.

Adukia, Anjali, Alex Eble, Emileigh Harrison, Hakizumwami Birali Runesha, and Teodora Szasz, "What We Teach About Race and Gender: Representation in Images and Text of Children's Books," *The Quarterly Journal of Economics*, 2023, *138* (4), 2225–2285.

_ , Richard Hornbeck, Daniel Keniston, and Benjamin Lualdi, "The Social Construction of Race during Reconstruction," *NBER Working Paper 33502*, 2025.

Agan, Amanda and Sonja Starr, "Ban the box, criminal records, and racial discrimination: A field experiment," *The Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

Aigner, Dennis J. and Glen G. Cain, "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review*, 1977, *30* (2), 175–187.

Ajzenman, Nicolás, Bruno Ferman, and Pedro C Sant'Anna, "Discrimination in the formation of academic networks: A field experiment on #EconTwitter," *American Economic Review: Insights*, 2025, *7* (3), 357–375.

Akerlof, George A and Rachel E Kranton, "Economics and Identity," *The Quarterly Journal of Economics*, 2000, *115*, 715–753.

Antman, Francisca and Brian Duncan, "Incentives to identify: Racial identity in the age of affirmative action," *Review of Economics and Statistics*, 2015, *97* (3), 710–713.

_ and _ , "American Indian Casinos and Native American Self-Identification," *Journal of the European Economic Association*, 2023, *21* (6), 2547–2585.

Arnold, David, Will Dobbie, and Crystal S. Yang, "Racial Bias in Bail Decisions," *Quarterly Journal of Economics*, 2018, *133* (4), 1885–1932.

_ , _ , and Peter Hull, "Measuring Racial Discrimination in Bail Decisions," *American Economic Review*, 2022, *112* (9), 2992–3038.

Arrow, Kenneth J., "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds., *Discrimination in Labor Markets*, Princeton, NJ: Princeton University Press, 1973, pp. 3–33.

Atkin, David, "The caloric costs of culture: Evidence from Indian migrants," *American Economic Review*, 2016, *106* (4), 1144–1181.

—, Eve Colson-Sihra, and Moses Shayo, "How do we choose our identity? A revealed preference approach using food consumption," *Journal of Political Economy*, 2021, *129* (4), 1193–1251.

Baerlocher, Diogo and Rodrigo Schneider, "Racial self-classification, group consciousness, and public employment representation," *Journal of Development Economics*, 2026, *181*, 103755.

Bailey, Stanley R., Mara Loveman, and Jeronimo O. Muniz, "Measures of "Race" and the analysis of racial inequality in Brazil," *Social Science Research*, 2013, *42* (1), 106–119.

Baron, E Jason, Joseph J Doyle Jr, Natalia Emanuel, Peter Hull, and Joseph Ryan, "Unwarranted Disparity in High-Stakes Decisions: Race Measurement and Policy Responses," in Randall Akee, Lawrence F. Katz, and Mark Loewenstein, eds., *Race, Ethnicity, and Economic Statistics for the 21st Century*, National Bureau of Economic Research, 2026.

Bayer, Patrick and Kerwin Kofi Charles, "Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940," *The Quarterly Journal of Economics*, 2018, *133* (3), 1459–1501.

—, —, and Ellora Derenoncourt, "Racial inequality in the labor market," in Christian Dustmann and Thomas Lemieux, eds., *Handbook of Labor Economics Vol. 6*, 2025, pp. 159–228.

Becker, Gary S., *The Economics of Discrimination*, Chicago: University of Chicago Press, 1957.

Ben-Ner, Avner, Brian P McCall, Massoud Stephane, and Hua Wang, "Identity and in-group/out-group differentiation in work and giving behaviors: Experimental evidence," *Journal of Economic Behavior & Organization*, 2009, *72* (1), 153–170.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *American Economic Review*, 2004, *94* (4), 991–1013.

Bingley, Paul and Alessandro Martinello, "Measurement Error in Income and Schooling, and the Bias for Linear Estimation," *Journal of Labor Economics*, 2017, *35* (4), 1117–1148.

Blinder, Alan S, "Wage Discrimination: Reduced Form and Structural Estimates," *Journal of Human resources*, 1973, *8* (4), 436–455.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg, "The Dynamics of Discrimination: Theory and Evidence," *American Economic Review*, 2019, *109* (10), 3395–3436.

___ , **Kareem Haggag, Alex Imas, and Devin G. Pope**, "Inaccurate Statistical Discrimination: An Identification Problem," *Review of Economics and Statistics*, 2025, *107* (3), 605–620.

___ , **Peter Hull, and Alex Imas**, "Systemic Discrimination: Theory and Measurement," *The Quarterly Journal of Economics*, 2025, *140* (3), 1743–1799.

Bueno, Natália S and Thad Dunning, "Race, Resources, and Representation: Evidence from Brazilian Politicians," *World Politics*, 2017, *69* (2), 327–365.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Cassan, Guilhem, "Identity-based policies and identity manipulation: Evidence from colonial Punjab," *American Economic Journal: Economic Policy*, 2015, *7* (4), 103–131.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," *American Economic Review*, 2024, *114* (5), 1488–1514.

Charness, Gary and Yan Chen, "Social Identity, Group Behavior, and Teams," *Annual Review of Economics*, 2020, *12*, 691–713.

Chor, Dóra, Alexandre Pereira, Antonio G. Pacheco, Ricardo V. Santos, Maria J. M. Fonseca, Maria I. Schmidt, Bruce B. Duncan, Sandhi M. Barreto, Estela M. L. Aquino, José G. Mill, Maria del C. B. Molina, Luana Giatti, Maria da C. Almeida, Isabela Bensenor, and Paulo A. Lotufo, "Context-dependence of race self-classification: Results from a highly mixed and unequal middle-income country," *PLOS ONE*, 2019, *14* (5), e0216653.

Cornwell, Christopher, Jason Rivera, and Ian M. Schmutte, "Wage Discrimination When Identity Is Subjective: Evidence from Changes in Employer-Reported Race," *Journal of Human Resources*, 2017, *52* (3), 719–755.

Dahis, Ricardo, Emily Nix, and Nancy Qian, "Choosing Racial Identity in the United States, 1880–1940," *NBER Working Paper 26465*, 2019. Revised June 2020.

Davenport, Lauren, "The Fluidity of Racial Classifications," *Annual Review of Political Science*, 2020, *23*, 221–240.

Davenport, Lauren D, Hakeem Jefferson, and Hunter E. Rendleman, "The Politics of Black Classification: Sociopolitical Cues and Racial Perception," 2026. Unpublished Manuscript.

de Holanda, Sérgio Buarque, *Raízes do Brasil*, Rio de Janeiro: Livraria José Olympio Editora, 1936.

de Lucena Coelho, Thiago, Fernanda Estevan, Marcos Nakaguma, and Alexandre Rabelo, "Do Black Politicians Matter? Political Leadership and Racial Composition in Top Public Sector Positions," *Political Leadership and Racial Composition in Top Public Sector Positions*, 2024.

Del Carpio, Lucia and Maria Guadalupe, "More women in tech? Evidence from a field experiment addressing social identity," *Management Science*, 2022, *68* (5), 3196–3218.

Dias, Felipe A., "How skin color, class status, and gender intersect in the labor market: Evidence from a field experiment," *Research in Social Stratification and Mobility*, 2020, *65*, 100477.

do Valle Silva, Nelson, "A research note on the cost of not being white in Brazil," *Studies in Comparative International Development*, 2000, *35*, 18–27.

Doleac, Jennifer L and Luke CD Stein, "The visible hand: Race and online market outcomes," *The Economic Journal*, 2013, *123* (572), F469–F492.

Evsyukova, Yulia, Felix Rusche, and Wladislaw Mill, "LinkedOut? A field experiment on discrimination in job network formation," *The Quarterly Journal of Economics*, 2025, *140* (1), 283–334.

Fernandes, Florestan, *A Integração do Negro na Sociedade de Classes*, São Paulo: Dominus Editora, 1965.

Ferraz, Claudio and Frederico Finan, "Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance," *Unpublished Manuscript*, 2011.

Firpo, Sergio, Michael França, and Alysson Portella, "Racial Inequality in the Brazilian Labor Market and the Role of Education," *Available at SSRN 3967828*, 2021.

Fournier-Montgieux, Alexandre, Hervé Le Borgne, Adrian Popescu, and Bertrand Luvison, "Reliable and Reproducible Demographic Inference for Fairness in Face Analysis," *arxiv preprint arXiv:2510.20482*, 2025.

França, Michael and Alysson Portella, *Números da discriminação racial: Desenvolvimento humano, equidade e políticas públicas*, Editora Jandaíra, 2023.

Francis, Andrew M and Maria Tannuri-Pianto, "Endogenous Race in Brazil : Affirmative Action and the Construction of Racial Identity among Young Adults," *Economic Development and Cultural Change*, 2013, *61*, 731–753.

Freeman, Jonathan B, Andrew M Penner, Aliya Saperstein, Matthias Scheutz, and Nalini Ambady, "Looking the part: Social status cues shape race perception," *PloS one*, 2011, *6* (9), e25107.

Freeman, Nicholas C., Edward E. Telles, and Rachel E. Goldberg, "The changing relationship between racial identity and skin color in Brazil," *Proceedings of the National Academy of Sciences*, 2025, *122* (1), e2411495121.

Freyre, Gilberto, *Casa-Grande & Senzala*, São Paulo: Companhia das Letras, 1933.

Fryer, Roland G, Devah Pager, and Jörg L Spenkuch, "Racial disparities in job finding and offered wages," *The Journal of Law and Economics*, 2013, *56* (3), 633–689.

G., Jr. Fryer Roland, "Racial Inequality in the 21st Century: The Declining Significance of Discrimination," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Elsevier, 2011, chapter 10, pp. 855–971.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative matching or exclusionary hiring? The impact of employment and pay policies on racial wage differences in Brazil," *American Economic Review*, 2021, *111* (10), 3418–3457.

Goldsmith, Arthur H., Darrick Hamilton, and Jr. Darity William A., "From Dark to Light: Skin Color and Wages Among African-Americans," *Journal of Human Resources*, 2007, *42* (4), 701–738.

_ , _ , and William Darity Jr., "Shades of Discrimination: Skin Tone and Wages," *American Economic Review*, 2006, *96* (2), 242–245.

Greiner, D. James and Donald B. Rubin, "Causal Effects of Perceived Immutable Characteristics," *Review of Economics and Statistics*, 2011, *93* (3), 775–785.

Hersch, Joni, "Skin-Tone Effects Among African Americans: Perceptions and Reality," *American Economic Review*, 2006, *96* (2), 251–255.

_ , "Profiling the New Immigrant Worker: The Effects of Skin Color and Height," *Journal of Labor Economics*, 2008, *26* (2), 345–386.

Holland, Paul W, "Causation and race," *White logic, white methods: Racism and methodology*, 2008, pp. 93–109.

Hu, Lily and Issa Kohler-Hausmann, "What is Perceived When Race is Perceived and Why It Matters for Causal Inference and Discrimination Studies," *Law & Society Review*, 2025, *59* (2), 239–264.

Hull, Peter, "One Weird Trick to Characterize Effective Populations in Design-Based Specifications," Technical Report, Metrics Note 2025.

IBGE, "Desigualdades Sociais por Cor ou Raça no Brasil," *Estudos e Pesquisas -Informação Demográfica e Socioeconômica*, 2022, *48* (2nd edition), 1–16.

Janusz, Andrew, "Electoral incentives and elite racial identification: Why Brazilian politicians change their race," *Electoral Studies*, 2021, *72*, 102340.

_ , "The electoral consequences of racial fluidity," *Electoral Studies*, 2023, *82*, 102597.

Jia, Ruixue and Torsten Persson, "Choosing ethnicity: The interplay between individual and social motives," *Journal of the European Economic Association*, 2021, *19* (2), 1203–1248.

Jr, Roland G Fryer and Steven D Levitt, "The causes and consequences of distinctively black names," *The Quarterly Journal of Economics*, 2004, *119* (3), 767–805.

Jr, William Darity, David K Guilkey, and William Winfrey, "Explaining differences in economic performance among racial and ethnic groups in the USA: the data examined," *American Journal of Economics and Sociology*, 1996, *55* (4), 411–425.

Kline, Patrick, Evan K Rose, and Christopher R Walters, "Systemic discrimination among large US employers," *The Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

_ , Raffaele Saggio, and Mikkel Sølvsten, "Leave-out estimation of variance components," *Econometrica*, 2020, *88* (5), 1859–1898.

Kohler-Hausmann, Issa, "Eddie Murphy and the Dangers of Counterfactual Causal Thinking about Detecting Racial Discrimination," *Northwestern University Law Review*, 2019, *113* (5), 1163–1227.

Koval, Christy Zhou and Ashleigh Shelby Rosette, "The Natural Hair Bias in Job Recruitment," *Social Psychological and Personality Science*, 2021, *12* (5), 741–750.

Kreisman, Daniel and Jonathan Smith, "Distinctively Black Names and Educational Outcomes," *Journal of Political Economy*, 2023, *131* (4), 877–897.

_ and Marcos A. Rangel, "On the Blurring of the Color Line: Wages and Employment for Black Males of Different Skin Tones," *The Review of Economics and Statistics*, 2015, *97* (1), 1–13.

Liebler, Carolyn A., Sonya R. Porter, Leticia E. Fernandez, James M. Noon, and Sharon R. Ennis, "America's Churning Races: Race and Ethnic Response Changes between Census 2000 and the 2010 Census," *Demography*, 2017, *54* (1), 259–284.

Lowe, Matt, "Types of contact: A field experiment on collaborative and adversarial caste integration," *American Economic Review*, 2021, *111* (6), 1807–1844.

Maggie, Yvonne and Peter Fry, "Qual Cor, Qual Raça?," *Estud. Av.*, 2004, *18* (50), 57–67.

Micheli, David De, "Racial reclassification and political identity formation," *World Politics*, 2021, *73* (1), 1–51.

Miller, Conrad and Ian M Schmutte, "The Dynamic Effects of Co-Racial Hiring," *Unpublished Manuscript*, 2023.

Miranda, Vítor, "A resurgence of black identity in Brazil? Evidence from an analysis of recent censuses," *Demographic Research*, 2015, *32*, 1603–1630.

Mitchell-Walthour, Gladys and William Darity Jr., "Choosing Blackness in Brazil's Racialized Democracy: The Endogeneity of Race in Salvador and São Paulo," *Latin American and Caribbean Ethnic Studies*, 2014, *9* (3), 318–348.

Mitrut, Andreea, Gabriel Kreindler, Margareta Matache, Andrei Munteanu, and Cristian Pop-Eleches, "Education and Selection into Ethnic Identification: Evidence from Roma People in Romania," *NBER Working Paper 34383*, 2025.

Monk, Ellis P., "Skin Tone Stratification among Black Americans, 2001–2003," *Social Forces*, 2014, *92* (4), 1313–1337.

_ , "The Consequences of "Race and Color" in Brazil," *Social Problems*, 2016, *63*, 413–430.

MTE, *Manual de Orientação da RAIS*, Brasília: Ministério do Trabalho e Emprego, 2024.

Muniz, Jerônimo and Stanley R. Bailey, "Does race response shift impact racial inequality?," *Demographic Research*, 2022, *47*, 935–966.

Muniz, Jerônimo, Aliya Saperstein, and Bernardo L. Queiroz, "Racial classification as a multistate process," *Demographic Research*, 2024, *50* (17), 457–472.

Neumark, David, "Experimental Research on Labor Market Discrimination," *Journal of Economic Literature*, 2018, *56* (3), 799–866.

Noghanibehambari, Hamid and Jason Fletcher, "Passing as White: Racial Identity and Old-Age Longevity," *NBER Working Paper 33394*, 2025.

Oaxaca, Ronald L., "Male-Female Wage Differentials in Urban Labor Markets," *International Economic Review*, 1973, *14* (3), 693–709.

Oh, Suanna, "Does identity affect labor supply?," *American Economic Review*, 2023, *113* (8), 2055–2083.

Osório, Rafael Guerreiro, "O sistema classificatório de cor ou raça do IBGE," Technical Report, Instituto de Pesquisa Econômica Aplicada (IPEA) 2003.

Pager, Devah, "The use of field experiments for studies of employment discrimination: Contributions, critiques, and directions for the future," *The Annals of the American Academy of Political and Social Science*, 2007, *609* (1), 104–133.

Penner, Andrew M and Aliya Saperstein, "How social status shapes race," *Proceedings of the National Academy of Sciences*, 2008, *105* (50), 19628–19630.

Pérez-Cervera, Laura, "Affirmative Action and Racial Voting," 2025. Unpublished Manuscript.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *American Economic Review*, 1972, *62* (4), 659–661.

Pickett, Robert EM, Aliya Saperstein, and Andrew M Penner, "Identification in interaction: Racial mirroring between interviewers and respondents," *Social Forces*, 2023, *102* (1), 23–44.

Quillian, Lincoln, Devah Pager, Ole Hexel, and Arnfinn H. Midtbøen, "Meta-analysis of Field Experiments Shows No Change in Racial Discrimination in Hiring over Time," *Proceedings of the National Academy of Sciences*, 2017, *114* (41), 10870–10875.

Rademakers, Robbert and André van Hoorn, "Ethnic switching: Longitudinal evidence on prevalence, correlates, and implications for measuring ethnic segregation," *Journal of Development Economics*, 2021, *152*, 102694.

Rodeheffer, Christopher D, Sarah E Hill, and Charles G Lord, "Does this recession make me look black? The effect of resource scarcity on the categorization of biracial faces," *Psychological Science*, 2012, *23* (12), 1476–1478.

Rose, Evan K, "A constructivist perspective on empirical discrimination research," *Journal of Economic Literature*, 2023, *61* (3), 906–923.

Saperstein, Aliya, "Capturing complexity in the United States: which aspects of race matter and when?," *Ethnic and Racial Studies*, 2012, *35* (8), 1484–1502.

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

Schroff, Florian, Dmitry Kalenichenko, and James Philbin, "FaceNet: A unified embedding for face recognition and clustering," in "Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition" 2015, pp. 815–823.

Schwarcz, Lilia Moritz, *Nem preto nem branco, muito pelo contrário: cor e raça na sociabilidade brasileira*, Editora Companhia das Letras, 2013.

Schwartzman, Luisa Farah, "Does money whiten? Intergenerational changes in racial classification in Brazil," *American Sociological Review*, 2007, *72* (6), 940–963.

Sen, Maya and Omar Wasow, "Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics," *Annual Review of Political Science*, 2016, *19*, 499–522.

Serengil, Sefik and Alper Ozpinar, "A Benchmark of Facial Recognition Pipelines and Co-Usability Performances of Modules," *Journal of Information Technologies*, 2024, *17* (2), 95–107.

Shayo, Moses, "Social identity and economic policy," *Annual Review of Economics*, 2020, *12*, 355–389.

Silveira, Leonardo, "Imputação da informação de raça/cor na Rais para o setor público brasileiro," Technical Report, Instituto de Pesquisa Econômica Aplicada (IPEA) 2022.

Small, Mario L and Devah Pager, "Sociological perspectives on racial discrimination," *Journal of Economic Perspectives*, 2020, *34* (2), 49–67.

Sorkin, Isaac, "Quantifying racial disparities using consecutive employment spells," Technical Report, National Bureau of Economic Research 2025.

Tajfel, Henri, *Human groups and social categories: Studies in social psychology*, Cambridge University Press, 1981.

_ **, Michael G. Billig, Robert P. Bundy, and Claude Flament**, "Social categorization and intergroup behaviour," *European Journal of Social Psychology*, 1971, *1* (2), 149–178.

Taylor, Shelley E., Susan T. Fiske, Nancy L. Etcoff, and Audrey J. Ruderman, "Categorical and contextual bases of person memory and stereotyping," *Journal of Personality and Social Psychology*, 1978, *36* (7), 778–793.

Telles, Edward E., *Race in Another America: The Significance of Skin Color in Brazil*, Princeton: Princeton University Press, 2004.

_ , *Pigmentocracies: Ethnicity, race, and color in Latin America*, UNC Press Books, 2014.

_ **and Nelson Lim**, "Does it matter who answers the race question? Racial classification and income inequality in Brazil," *Demography*, 1998, *35* (4), 465–474.

_ **and Tianna S. Paschel**, "Who Is Black, White, or Mixed Race? How Skin Color, Status, and Nation Shape Racial Classification in Latin America," *American Journal of Sociology*, 2014, *120* (3), 864–907.

Woo-Mora, L. Guillermo, "Unveiling the Cosmic Race: Skin tone and intergenerational economic disparities in Latin America and the Caribbean," *Journal of Development Economics*, 2026, *179*, 103594.

Online Appendix to
"Perceptions of Race in the Labor Market"

Pedro C. Sant'Anna Sulin Sardoschau Aiko Schmeißer

23rd May 2026

A Additional Figures

Figure A.1: Histogram of Non-White Share Across Establishments

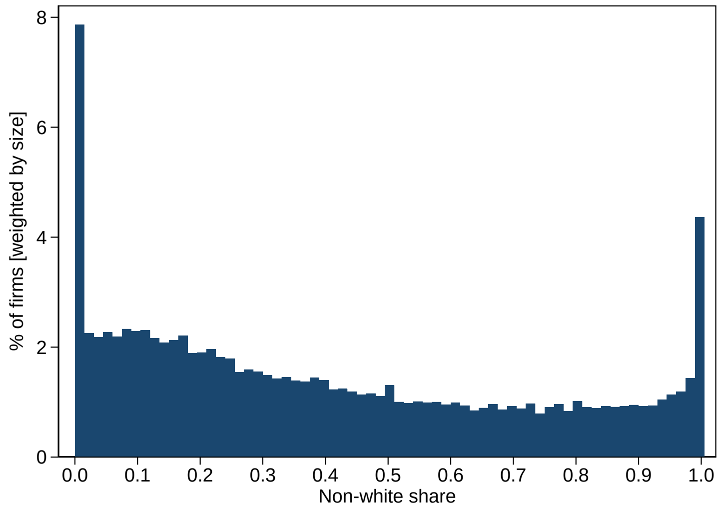

Notes: The figure displays a histogram of the non-White (Black or Brown) share of workers across establishments in the full RAIS sample. Establishments are weighted by their number of workers.

Figure A.2: Skin Tone Distribution, Detailed Racial Categories

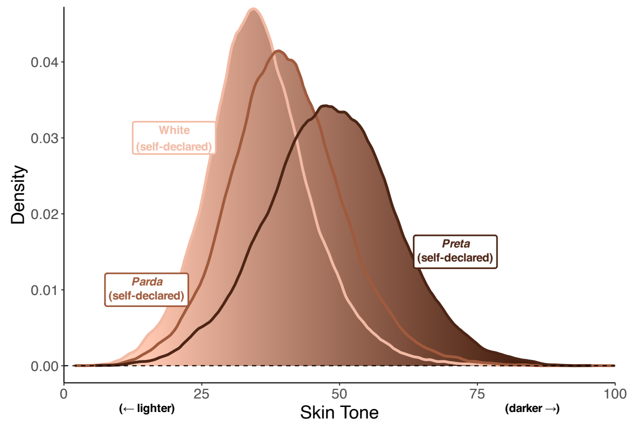

Notes: The figure shows the distribution of skin tone as measured from the electoral photos of candidates, splitting candidates by their self-declared race, distinguishing White, *Pardo* (Brown), and *Preto* (Black) candidates.

Figure A.3: Skin Tone Distribution by Mismatch Category

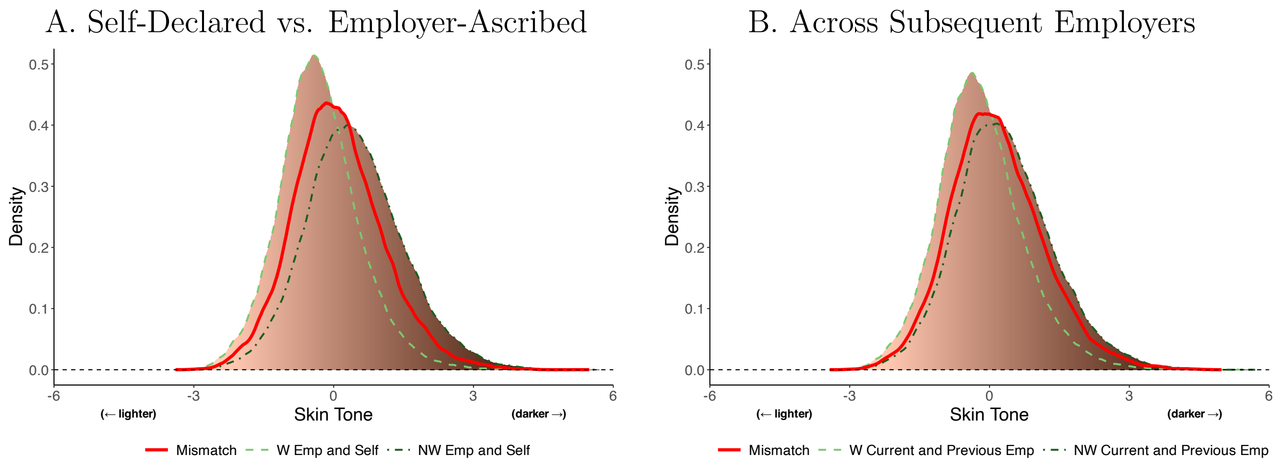

Notes: The figures show histograms of skin tone by categories of racial mismatch. The samples are the same as in Table 1. Panel A shows the skin tone distribution by cells of self-declared and employer-ascribed race, while Panel B shows the skin tone distribution by cells of ascribed race by the current and previous employer. The skin tone measure is standardized to have a mean of zero and a standard deviation of one within our sample.

Figure A.4: Wage Disparities, Heterogeneity by Gender

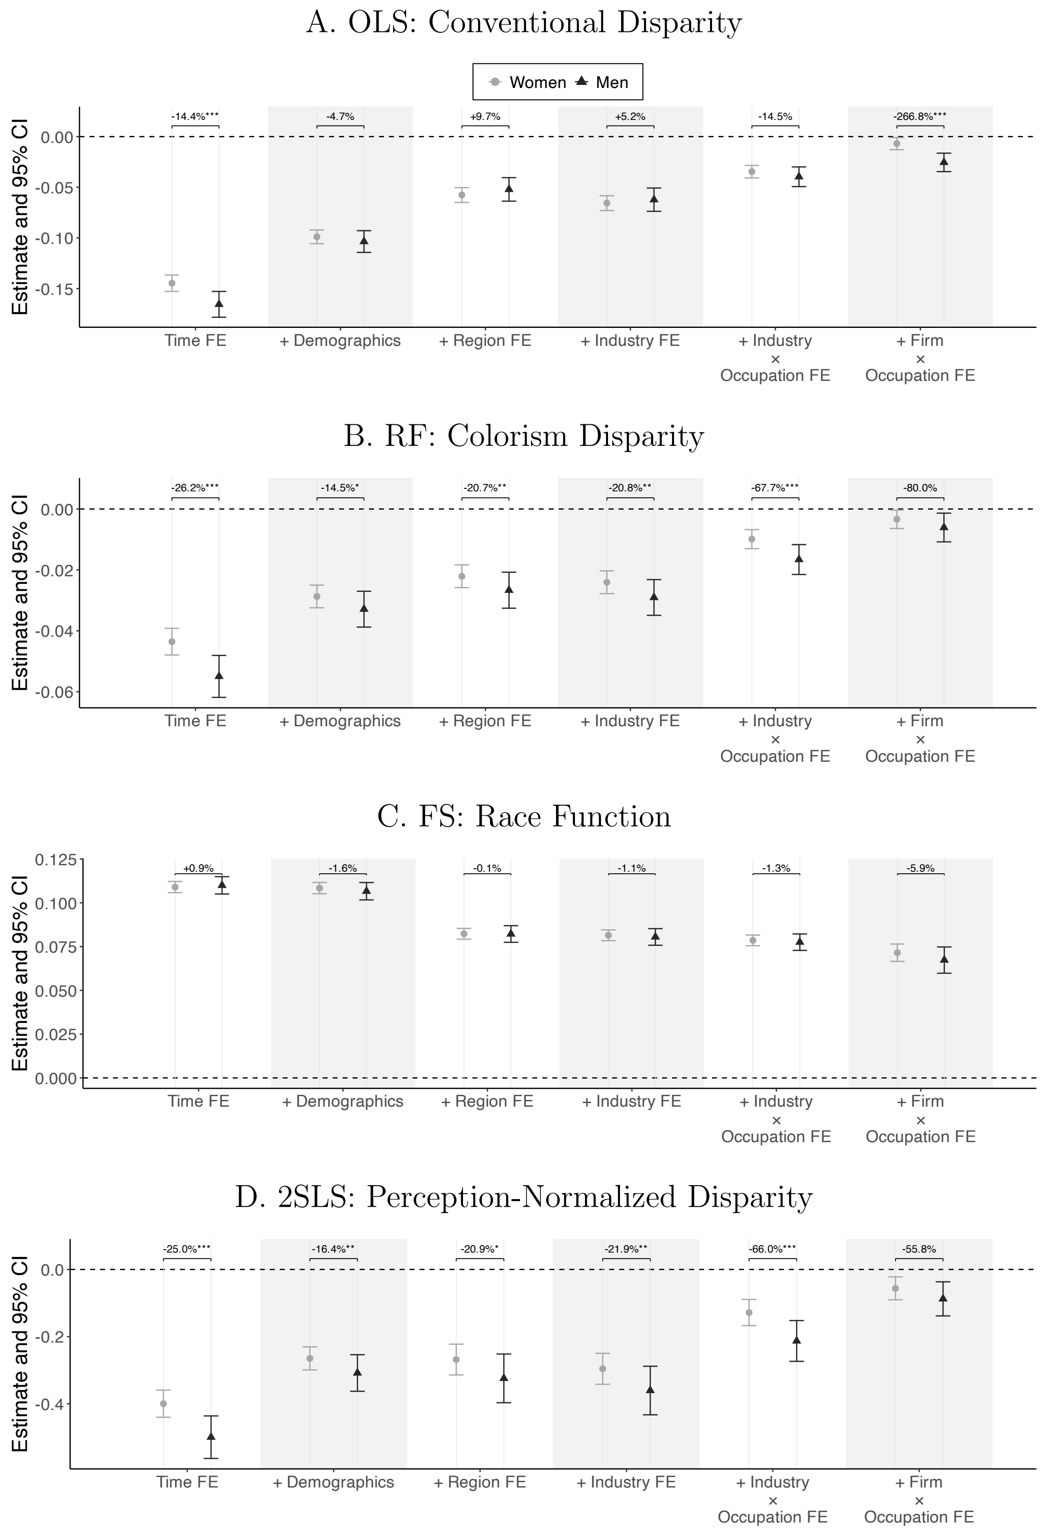

Notes: The figures display different estimates of racial wage gaps for female and male workers. We use the same specifications as columns (1) to (6) of Table 3, with interaction terms for each gender. present point estimates and 95% confidence intervals computed using clustered standard errors at the individual level. The brackets above pairs of point estimates report the percent difference between the estimates (relative to the absolute value of the reference category), with stars indicating whether the difference is statistically significant. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A.5: Wage Disparities, Heterogeneity by Median Age

A. OLS: Conventional Disparity

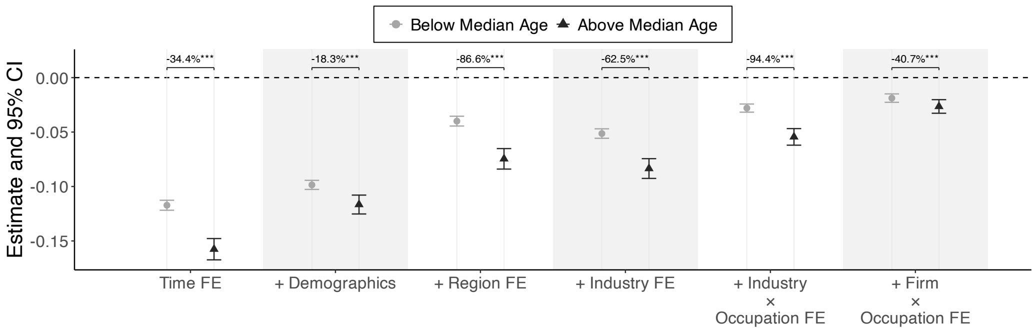

B. RF: Colorism Disparity

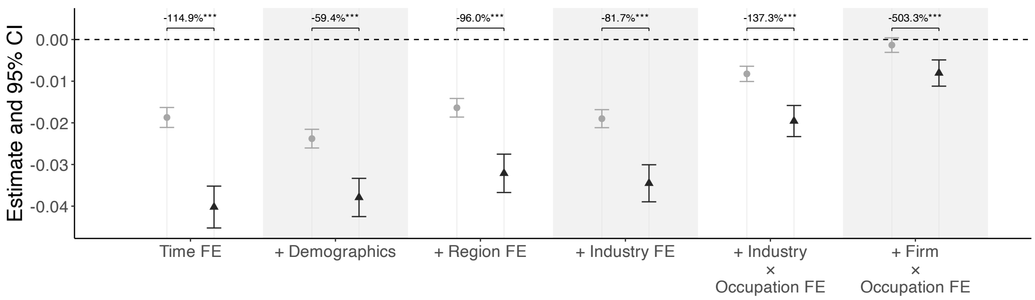

C. FS: Race Function

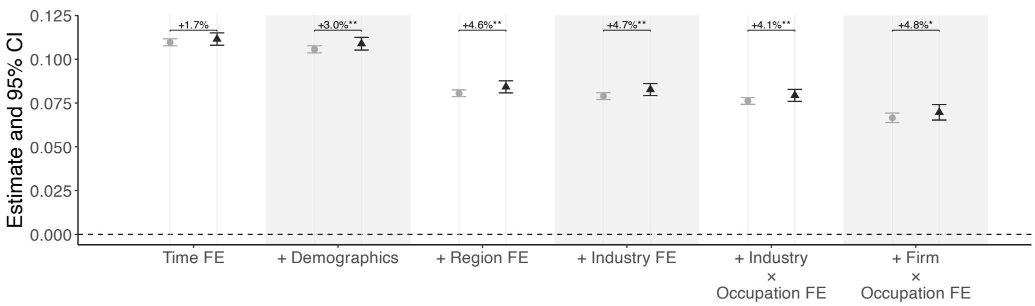

D. 2SLS: Perception-Normalized Disparity

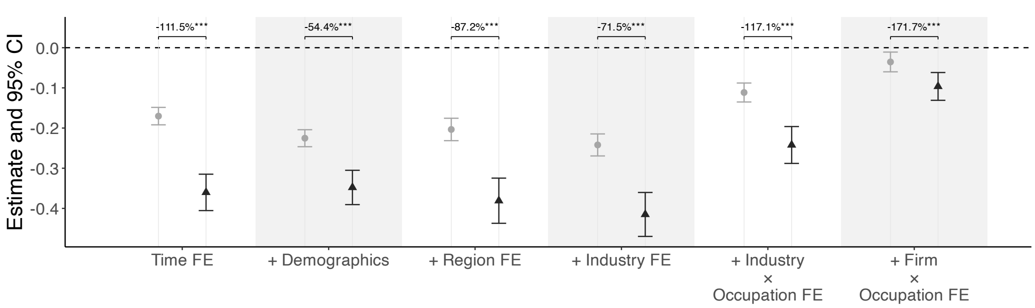

Notes: The figures display different estimates of racial wage gaps for workers above or below the median age in the sample. We use the same specifications as columns (1) to (6) of Table 3, with interaction terms for being above/below the median age. present point estimates and 95% confidence intervals computed using clustered standard errors at the individual level. The brackets above pairs of point estimates report the percent difference between the estimates (relative to the absolute value of the reference category), with stars indicating whether the difference is statistically significant. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Additional Tables

Table B.1: Summary Statistics of Candidate Sample and Full RAIS Sample

	(1) Candidates Mean	(2) Non-candidates Mean	(3) Difference (1) - (2)
Worker Characteristics			
Non-White (Employer-Ascribed)	0.37	0.37	-0.00**
Male	0.75	0.62	0.13***
Age (years)	34.48	33.65	0.84***
Log Hourly Wage (2018 BRL)	2.18	2.20	-0.02***
Highest Degree Attained			
Less than Elementary	0.03	0.04	-0.01***
Elementary School	0.14	0.15	-0.01***
Middle School	0.23	0.23	0.00***
High School	0.52	0.49	0.03***
Undergraduate or more	0.09	0.10	-0.01***
Establishment Characteristics			
Firm Size	536.54	580.17	-43.62***
Average Log Hourly Wage	2.15	2.21	-0.06***
Sector			
Agriculture	0.03	0.02	0.01***
Extractive industry	0.01	0.01	0.00***
Manufacturing	0.24	0.25	-0.01***
Construction	0.06	0.07	-0.00***
Retail	0.25	0.25	-0.01***
Transport & communication	0.13	0.13	-0.00***
Finance & real state	0.15	0.17	-0.02***
Education & health	0.09	0.07	0.01***
Other services	0.05	0.03	0.02***
Region			
North	0.05	0.04	0.01***
Northeast	0.17	0.16	0.01***
Center-West	0.51	0.55	-0.04***
Southeast	0.19	0.18	0.02***
South	0.08	0.07	0.01***
No. of worker-year obs.	1,510,833	310,051,787	
No. of workers	331,531	61,075,388	
No. of establishments	358,899	5,616,123	

Notes: The table shows mean characteristics of workers who have run for a political office between 2014 and 2024 (our estimation sample, column (1)) and all other workers in RAIS (column (2)). Column (3) shows the difference in means between both samples. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.2: Racial Gaps in Candidate Sample vs Full RAIS Sample

	(1)	(2)	(3)	(4)	(5)	(6)
		Candidates			Non-candidates	
	White Mean	Non-White Mean	Racial gap (1)-(2)	White Mean	Non-White Mean	Racial gap (4)-(5)
Worker Characteristics						
Male	0.73	0.77	-0.04***	0.60	0.66	-0.06***
Age (years)	34.72	34.09	0.63***	33.87	33.25	0.62***
Log Hourly Wage (2018 BRL)	2.22	2.10	0.13***	2.28	2.06	0.23***
Highest Degree Attained						
Less than Elementary	0.02	0.04	-0.02***	0.03	0.06	-0.03***
Elementary School	0.13	0.15	-0.02***	0.14	0.16	-0.03***
Middle School	0.23	0.23	0.00	0.23	0.24	-0.01***
High School	0.51	0.52	-0.01***	0.49	0.49	-0.01***
Undergraduate or more	0.11	0.06	0.05***	0.12	0.05	0.07***
Establishment Characteristics						
Firm Size	441.86	698.04	-256.17***	501.76	707.93	-206.17***
Average Log Hourly Wage	2.18	2.09	0.09***	2.27	2.10	0.17***
Sector						
Agriculture	0.02	0.03	-0.01***	0.01	0.02	-0.01***
Extractive industry	0.01	0.01	-0.00***	0.01	0.01	-0.00***
Manufacturing	0.26	0.21	0.05***	0.27	0.22	0.05***
Construction	0.05	0.09	-0.05***	0.05	0.10	-0.05***
Retail	0.26	0.23	0.03***	0.26	0.25	0.01***
Transport & communication	0.13	0.13	-0.00***	0.13	0.13	0.00
Finance & real state	0.13	0.17	-0.04***	0.16	0.18	-0.02***
Education & health	0.09	0.08	0.01***	0.08	0.07	0.01***
Other services	0.05	0.05	0.01***	0.03	0.03	0.00***
Region						
North	0.02	0.11	-0.08***	0.02	0.09	-0.07***
Northeast	0.09	0.31	-0.22***	0.08	0.30	-0.22***
Southeast	0.55	0.44	0.12***	0.60	0.48	0.12***
South	0.28	0.05	0.23***	0.25	0.04	0.21***
Center-West	0.06	0.10	-0.04***	0.06	0.09	-0.04***
N	952,450	558,383		192,996,780	113,674,338	

Notes: The table shows racial gaps (by employer-ascribed race) in mean characteristics of workers who have run for a political office between 2014 and 2024 (our estimation sample, columns (1)-(3)) and all other workers in RAIS (columns (4)-(6)). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.3: Summary Statistics of Candidate Sample by Detailed Racial Group

	(1) White Mean	(2) *Parda* Mean	(3) *Preta* Mean	(4) Difference (1)-(2)	(5) Difference (2)-(3)
Worker Characteristics					
Gender (1=male)	0.72	0.78	0.77	-0.06***	0.01***
Age (years)	34.83	33.89	34.62	0.94***	-0.72***
Log Hourly Wage (2018 BRL)	2.25	2.09	2.08	0.16***	0.01***
Highest degree attained					
Less than Elementary	0.02	0.04	0.04	-0.02***	-0.00***
Elementary School	0.12	0.15	0.17	-0.03***	-0.02***
Middle School	0.22	0.25	0.26	-0.02***	-0.01***
High School	0.52	0.51	0.49	0.01***	0.02***
Undergraduate or more	0.12	0.05	0.05	0.06***	0.01***
Establishment Characteristics					
Firm Size	478.33	597.91	638.79	-119.57***	-40.89***
Average Log Hourly Wage	2.19	2.09	2.12	0.11***	-0.03***
Sector					
Agriculture	0.02	0.03	0.03	-0.01***	-0.00
Extractive industry	0.01	0.01	0.01	-0.00***	0.00
Manufacturing	0.26	0.22	0.23	0.05***	-0.01***
Construction	0.05	0.09	0.08	-0.04***	0.00***
Retail	0.26	0.24	0.21	0.01***	0.03***
Transport & communication	0.12	0.13	0.13	-0.01***	0.01***
Finance & real state	0.14	0.16	0.17	-0.02***	-0.01***
Education & health	0.09	0.07	0.08	0.02***	-0.01***
Other services	0.05	0.05	0.06	0.01***	-0.01***
Region					
North	0.02	0.11	0.04	-0.08***	0.06***
Northeast	0.09	0.28	0.21	-0.19***	0.07***
Southeast	0.54	0.44	0.57	0.10***	-0.13***
South	0.29	0.07	0.11	0.22***	-0.04***
Center-West	0.06	0.11	0.07	-0.05***	0.04***
No. of worker-year obs.	829,603	522,553	158,677	1,352,156	681,230
No. of workers	176,606	120,426	34,499		

Notes: The table shows mean characteristics of workers who have run for a political office between 2014 and 2024, split by their self-declared race (White, *Parda*, and *Preta*). Columns (4) and (5) show the difference in means between workers in the White vs. *Parda*, and *Parda* vs. *Preta* groups, respectively. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.4: Mismatch in Ascribed Race, Detailed Racial Categories

Panel A: Mismatch between Self-Declared and Employer-Ascribed Race

	Employer-Ascribed Race					
	White		*Parda*		*Preta*	
Self-Declared Race	%	Skin Tone	%	Skin Tone	%	Skin Tone
White	45.20%	-0.32	9.22%	-0.13	0.49%	0.14
Parda	14.86%	0.07	17.82%	0.21	1.90%	0.55
Preta	2.99%	0.77	4.16%	0.80	3.35%	1.19

Panel B: Mismatch across Subsequent Employers

	Employer-Ascribed Race, Current					
	White		*Parda*		*Preta*	
Employer-Ascribed Race, Previous	%	Skin Tone	%	Skin Tone	%	Skin Tone
White	49.82%	-0.22	10.79%	0.08	1.38%	0.71
Parda	9.57%	0.07	20.71%	0.24	1.73%	0.81
Preta	1.31%	0.72	1.83%	0.75	2.87%	1.11

Notes: The table displays the (mis)match shares for different race measures in our sample, as well as the average skin tone for each cell. This table considered the detailed racial categories available in Brazilian statistics, White, *Parda* (brown), and *Preta* (black), excluding indigenous and Asian individuals. Skin tone is standardized to have a mean of zero and a standard deviation of one across the whole sample (higher values represent darker shades, while lower values represent lighter shades). Panel A presents the distribution of workers by self-declared race (rows) and employer-ascribed race (columns). It considers all worker-year observations ($N = 1,510,833$). Panel B presents the distribution of ascribed race for the same worker across different employers in consecutive years. It considers all workers who changed jobs ($N = 153,366$) and presents the distribution of workers by their ascribed race in the previous employer (rows) and current employer (columns).

Table B.5: Mismatch in Ascribed Race - Robustness

[A] Mismatch between Self-Declared and Employer-Ascribed Race

Employer-ascribed Self-declared	White White	Non-White Non-White	White Non-White	Non-White White	*N*
Baseline sample	45.20%	27.25%	17.84%	9.71%	1,510,833
Firms with variation in ascribed race	42.30%	29.24%	18.08%	10.38%	1,328,249
Firms with HR employees	43.48%	28.88%	17.24%	10.41%	910,506
Self-declared race in same period	43.28%	29.23%	16.04%	11.45%	48,316
Councilor candidates	44.97%	27.49%	17.99%	9.55%	1,412,050

Employer-ascribed Self-declared	Male Male	Female Female	Male Female	Female Male	*N*
Baseline sample	73.75%	24.26%	1.06%	0.93%	1,510,833

[B] Mismatch across Subsequent Employers

Employer-ascribed, current Employer-ascribed, previous	White White	Non-White Non-White	White Non-White	Non-White White	*N*
Baseline sample	49.82%	27.12%	10.88%	12.18%	153,366
Firms with variation in ascribed race	45.68%	29.96%	11.51%	12.85%	128,142
Firms with HR employees	47.74%	30.64%	10.22%	11.40%	71,501
All workers in RAIS	50.44%	26.70%	10.78%	12.08%	33,898,971

Employer-ascribed, current Employer-ascribed, previous	Male Male	Female Female	Male Female	Female Male	*N*
Baseline sample	76.58%	20.76%	1.32%	1.34%	153,366

Notes: The table shows mismatch rates between self-declared and employer-ascribed race (Panel A) and in employer-ascribed race across different employers in subsequent years (Panel B) in alternative subsamples. Specifically, we restrict the sample to firms that do not ascribe the same racial category to all their workers, firms that have at least one employee in an HR occupation, jobs in 2015 and candidates running in 2014 or 2016 (such that self-declared and employer-ascribed race are report at roughly the same time), as well as candidates that run for local councilor positions. Panel B also considers job changes among all workers in RAIS (not limited to the candidate sample). Both panels also report mismatch rates across reported gender instead of race.

Table B.6: Mismatch in Ascribed Race - Aggregating Across Employers

	(1) OLS	(2) OLS	(3) SSIV	(4) OLS	(5) SSIV
Non-White Employer	0.454***				
	(0.002)				
Non-White Employer Mode		0.472***	0.613***		
		(0.002)	(0.003)		
Non-White Employer Mean				0.597***	0.659***
				(0.002)	(0.002)
Constant	0.283***	0.286***	0.240***	0.230***	0.207***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
N	1,510,833	1,510,833	1,439,433	1,510,833	1,439,433

Notes: The table reports results from regressions of self-declared race (indicator equal to one if the worker identifies as non-White) on employer-ascribed race (indicator equal to one if the worker is ascribed non-White). While column (1) uses race as ascribed by the current employer, columns (2) to (5) aggregate across employers by considering the mode or mean across all yearly observations of the worker. Columns (3) and (5) implement a split-sample IV (SSIV) approach, in which the mode or mean constructed from one randomly selected half of a worker's yearly observations is instrumented with the mode or mean constructed from the other half of observations. Standard errors clustered at the individual level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.7: Determinants of Employer-Ascribed Race including Employment History Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Skin Tone	0.110***	0.104***	0.082***	0.081***	0.079***	0.068***	0.040***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Non-White Self-Declared							0.244***
							(0.007)
Demographics							
Male		0.037***	0.016***	0.009**	0.001	-0.000	-0.002
		(0.004)	(0.004)	(0.004)	(0.004)	(0.009)	(0.008)
Age		0.007***	0.003*	0.001	-0.000	-0.002	-0.002
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Age^2		-0.010***	-0.005***	-0.004***	-0.003**	-0.001	0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
(Some) High School		-0.001	-0.028***	-0.019***	-0.010**	-0.008	-0.003
		(0.004)	(0.004)	(0.004)	(0.004)	(0.007)	(0.007)
(Some) University		-0.044***	-0.087***	-0.078***	-0.040***	-0.014	0.000
		(0.006)	(0.005)	(0.006)	(0.006)	(0.017)	(0.015)
Work History							
Tenure		-0.055***	-0.043***	-0.036***	-0.035***	-0.048***	-0.045***
		(0.003)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)
Tenure^2		0.005***	0.005***	0.003**	0.003*	0.002**	0.002*
		(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Final Wage in Previous Job		-0.046***	-0.030***	-0.033***	-0.028***	-0.015***	-0.010**
		(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)
Region							
Northeast			-0.057***	-0.060***	-0.062***		
			(0.008)	(0.008)	(0.008)		
Southeast			-0.374***	-0.370***	-0.372***		
			(0.007)	(0.007)	(0.007)		
South			-0.569***	-0.560***	-0.561***		
			(0.007)	(0.007)	(0.007)		
Center-West			-0.191***	-0.188***	-0.188***		
			(0.009)	(0.009)	(0.009)		
Constant	0.370***	0.340***	0.713***	0.710***	0.707***	0.369***	0.258***
	(0.001)	(0.005)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)
N	1,510,833	374,036	374,036	374,006	373,905	298,044	298,044
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes
R-squared	0.057	0.065	0.209	0.213	0.223	0.779	0.793
Adjusted R-squared	0.057	0.065	0.209	0.213	0.222	0.719	0.736

Notes: The table reports OLS estimates of the relationship between employer-ascribed race (an indicator equal to one for non-White) and various covariates, enriching the analysis of Table 2 with work history variables. The sample in columns 2-7 only includes workers whom we observe in multiple establishments. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.8: Determinants of Self-Declared Race

	(1)	(2)	(3)	(4)	(5)	(6)
Skin Tone	0.154***	0.149***	0.126***	0.125***	0.122***	0.115***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Demographics						
Male		0.007***	0.002	-0.001	-0.005**	-0.006
		(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Age		-0.025***	-0.017***	-0.019***	-0.014***	-0.009***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Age^2		-0.011***	-0.008***	-0.007***	-0.007***	-0.006***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
(Some) High School		-0.033***	-0.054***	-0.050***	-0.031***	-0.025***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
(Some) University		-0.159***	-0.173***	-0.174***	-0.107***	-0.077***
		(0.003)	(0.003)	(0.003)	(0.004)	(0.008)
Region						
Northeast			-0.071***	-0.072***	-0.072***	
			(0.004)	(0.004)	(0.004)	
Southeast			-0.344***	-0.342***	-0.343***	
			(0.004)	(0.004)	(0.004)	
South			-0.535***	-0.529***	-0.531***	
			(0.004)	(0.004)	(0.004)	
Center-West			-0.180***	-0.179***	-0.177***	
			(0.005)	(0.005)	(0.005)	
Constant	0.451***	0.496***	0.815***	0.813***	0.796***	0.483***
	(0.001)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
N	1,510,833	1,510,833	1,510,833	1,510,510	1,510,445	1,249,942
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No
Industry $\times$ Occupation FE	No	No	No	No	Yes	No
Firm $\times$ Occupation FE	No	No	No	No	No	Yes
R-squared	0.097	0.110	0.214	0.217	0.226	0.772
Adjusted R-squared	0.097	0.110	0.214	0.217	0.225	0.711

Notes: The table reports OLS estimates of the relationship between self-declared race (an indicator equal to one for non-White) and various covariates. The first column displays the relationship with skin tone, when only controlling for year fixed effects. The subsequent columns progressively add demographic covariates as well as region, one-digit industry, three-digit occupation, and firm fixed effects. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.9: Correlates of Self-Declared Race including Employment History Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Skin Tone	0.154***	0.144***	0.125***	0.124***	0.122***	0.115***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Demographics						
Male		0.046***	0.028***	0.024***	0.011**	0.009
		(0.004)	(0.004)	(0.004)	(0.005)	(0.010)
Age		0.002	-0.001	-0.003	-0.004**	-0.002
		(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Age^2		-0.013***	-0.009***	-0.008***	-0.007***	-0.006***
		(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
(Some) High School		-0.019***	-0.042***	-0.037***	-0.025***	-0.021**
		(0.004)	(0.004)	(0.004)	(0.004)	(0.009)
(Some) University		-0.104***	-0.140***	-0.136***	-0.090***	-0.058***
		(0.007)	(0.006)	(0.006)	(0.007)	(0.019)
Work History						
Tenure		-0.037***	-0.026***	-0.021***	-0.019***	-0.012***
		(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Tenure^2		0.007***	0.007***	0.006***	0.005***	0.002**
		(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Final Wage in Previous Job		-0.046***	-0.033***	-0.035***	-0.027***	-0.021***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.006)
Region						
Northeast			-0.067***	-0.070***	-0.072***	
			(0.009)	(0.009)	(0.009)	
Southeast			-0.324***	-0.321***	-0.325***	
			(0.008)	(0.008)	(0.008)	
South			-0.515***	-0.507***	-0.510***	
			(0.008)	(0.008)	(0.008)	
North			-0.160***	-0.157***	-0.157***	
			(0.010)	(0.010)	(0.010)	
Constant	0.451***	0.433***	0.763***	0.760***	0.759***	0.454***
	(0.001)	(0.006)	(0.009)	(0.009)	(0.009)	(0.011)
N	1,510,833	374,036	374,036	374,006	373,905	298,044
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No
Industry × Occupation FE	No	No	No	No	Yes	No
Firm × Occupation FE	No	No	No	No	No	Yes
R-squared	0.097	0.110	0.214	0.217	0.226	0.772
Adjusted R-squared	0.097	0.110	0.214	0.217	0.225	0.711

Notes: The table reports regression estimates of the relationship between self-declared race (an indicator equal to one for non-White) and various covariates, enriching the analysis of Table B.8 with work history variables. The sample in columns 2-7 only includes workers whom we observe in multiple establishments. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.10: Determinants of Employer-Ascribed Race with Machine-Learning Prediction of Visual Cues

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
P(Non-White Employer-Ascribed)	1.030***	1.014***	0.828***	0.820***	0.806***	0.676***	0.492***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.007)	(0.008)
Non-White Self-Declared							0.163***
							(0.003)
Demographics							
Male		0.007***	0.001	-0.007***	-0.009***	-0.011***	-0.011***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Age		-0.005***	0.001	-0.000	0.000	-0.000	-0.002
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age^2		-0.004***	-0.001**	-0.001	-0.001	-0.000	0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
(Some) High School		0.003*	-0.021***	-0.015***	-0.007***	-0.004	-0.001
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
(Some) University		-0.044***	-0.070***	-0.064***	-0.034***	-0.018***	-0.009
		(0.003)	(0.002)	(0.003)	(0.003)	(0.006)	(0.006)
Region							
Northeast			-0.051***	-0.052***	-0.054***		
			(0.004)	(0.004)	(0.004)		
Southeast			-0.368***	-0.366***	-0.367***		
			(0.003)	(0.003)	(0.003)		
South			-0.506***	-0.501***	-0.501***		
			(0.004)	(0.004)	(0.004)		
Center-West			-0.193***	-0.191***	-0.191***		
			(0.004)	(0.004)	(0.004)		
Constant	-0.008***	0.001	0.396***	0.398***	0.397***	0.135***	0.125***
	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
N	1,409,846	1,409,846	1,409,846	1,409,546	1,409,476	1,164,572	1,164,572
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes
R-squared	0.167	0.168	0.276	0.279	0.286	0.796	0.802
Adjusted R-squared	0.167	0.168	0.276	0.279	0.286	0.741	0.747

Notes: The table reports OLS estimates of the relationship between employer-ascribed race (an indicator equal to one for non-White) and various covariates, conditioning on a machine-learning prediction of employer ascriptions based on candidate photos. See Appendix C.3 for details on the prediction algorithm. The first column displays the coefficient of the predicted probability to be non-White, when only controlling for year fixed effects. The subsequent columns progressively add demographic covariates as well as region, one-digit industry, three-digit occupation, and firm fixed effects. The last column also conditions on individuals' self-declared race. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.11: Determinants of Employer-Ascribed Race with Machine-Learning Prediction of Visual Cues, including Employment History Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
P(Non-White Employer-Ascribed)	1.026***	0.982***	0.826***	0.818***	0.809***	0.680***	0.482***
	(0.004)	(0.008)	(0.009)	(0.009)	(0.009)	(0.016)	(0.018)
Non-White Self-Declared							0.167***
							(0.008)
Demographics							
Male		0.025***	0.008**	0.001	-0.004	-0.006	-0.005
		(0.004)	(0.004)	(0.004)	(0.004)	(0.009)	(0.008)
Age		0.019***	0.014***	0.012***	0.011***	0.008**	0.005
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Age^2		-0.007***	-0.003*	-0.002	-0.001	-0.000	0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
(Some) High School		0.007*	-0.018***	-0.010***	-0.005	-0.001	0.001
		(0.004)	(0.004)	(0.004)	(0.004)	(0.007)	(0.007)
(Some) University		-0.008	-0.052***	-0.043***	-0.020***	-0.001	0.005
		(0.006)	(0.005)	(0.005)	(0.006)	(0.016)	(0.015)
Work History							
Tenure		-0.047***	-0.037***	-0.031***	-0.031***	-0.043***	-0.042***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)
Tenure^2		0.004***	0.004***	0.003*	0.003*	0.002**	0.002*
		(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Final Wage in Previous Job		-0.033***	-0.020***	-0.023***	-0.021***	-0.010**	-0.008*
		(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)
Region							
Northeast			-0.063***	-0.065***	-0.068***		
			(0.008)	(0.008)	(0.008)		
Southeast			-0.356***	-0.352***	-0.353***		
			(0.008)	(0.008)	(0.008)		
South			-0.501***	-0.494***	-0.494***		
			(0.008)	(0.008)	(0.008)		
Center-West			-0.170***	-0.168***	-0.168***		
			(0.009)	(0.009)	(0.009)		
Constant	-0.006***	-0.023***	0.380***	0.380***	0.382***	0.116***	0.113***
	(0.002)	(0.005)	(0.009)	(0.009)	(0.009)	(0.011)	(0.010)
N	1,409,846	348,093	348,093	348,065	347,958	277,006	277,006
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes
R-squared	0.166	0.175	0.270	0.274	0.285	0.839	0.843
Adjusted R-squared	0.166	0.175	0.270	0.274	0.283	0.787	0.793

Notes: The table reports OLS estimates of the relationship between employer-ascribed race (an indicator equal to one for non-White) and various covariates, conditioning on a machine-learning prediction of employer ascriptions based on candidate photos and enriching the analysis of Table B.10 with work history variables. The sample in columns 2-7 only includes workers whom we observe in multiple establishments. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.12: Wage Disparities, Employer-Ascribed Race, Only Candidates to City Council

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
[A] OLS: Conventional Disparity							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Non-White Employer-Ascribed	-0.131***	-0.100***	-0.046***	-0.056***	-0.035***	-0.021***	-0.018***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Non-White Self-Declared							-0.013***
							(0.002)
[B] Reduced Form: Colorism Disparity							
Skin Tone	-0.018***	-0.031***	-0.024***	-0.026***	-0.015***	-0.005***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							-0.015***
							(0.002)
[C] First Stage: Race Function							
Skin Tone	0.109***	0.106***	0.081***	0.080***	0.077***	0.068***	0.041***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							0.237***
							(0.003)
[D] 2SLS: Perception-Normalized Disparity							
Non-White Employer-Ascribed	-0.169***	-0.290***	-0.293***	-0.325***	-0.196***	-0.079***	-0.090***
	(0.011)	(0.011)	(0.014)	(0.014)	(0.012)	(0.013)	(0.022)
Non-White Self-Declared							0.006
							(0.006)
N	1,329,234	1,329,234	1,329,234	1,328,946	1,328,882	1,093,180	1,093,180
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	Yes	Yes	Yes	No	No
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes

Notes: The table reports different measures of racial wage gaps based on employer-ascribed race for the subsample of candidates who exclusively ran for city council in our sample. All columns include year fixed-effects. The first column always shows the unconditional relationship (apart from year fixed-effects), and the remaining columns progressively add controls. The final column includes all fixed-effects, as well as controls for self-declared race. Panel A shows coefficient estimates of OLS regressions of log hourly wage on employer-ascribed race. Panel B shows reduced form regressions of log hourly wage on skin tone. Panel C presents the first-stage coefficient from regressing employer-ascribed race on skin tone. Finally, Panel D shows 2SLS estimates of regressing log hourly wage on employer-ascribed race, instrumented by skin tone. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.13: Wage Disparities, Employer-Ascribed Race, Only Candidates Before Running for Election

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
[A] OLS: Conventional Disparity							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Non-White Employer-Ascribed	-0.147***	-0.105***	-0.054***	-0.063***	-0.039***	-0.022***	-0.018***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Non-White Self-Declared							-0.015***
							(0.002)
[B] Reduced Form: Colorism Disparity							
Skin Tone	-0.023***	-0.031***	-0.025***	-0.027***	-0.015***	-0.005***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							-0.018***
							(0.002)
[C] First Stage: Race Function							
Skin Tone	0.106***	0.103***	0.079***	0.078***	0.075***	0.066***	0.040***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-White Self-Declared							0.237***
							(0.003)
[D] 2SLS: Perception-Normalized Disparity							
Non-White Employer-Ascribed	-0.212***	-0.304***	-0.313***	-0.346***	-0.198***	-0.074***	-0.072***
	(0.012)	(0.011)	(0.015)	(0.014)	(0.012)	(0.013)	(0.023)
Non-White Self-Declared							-0.001
							(0.006)
N	1,346,003	1,346,003	1,346,003	1,345,690	1,345,621	1,097,213	1,097,213
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	Yes	Yes	Yes	No	No
Industry FE	No	No	No	Yes	No	No	No
Industry × Occupation FE	No	No	No	No	Yes	No	No
Firm × Occupation FE	No	No	No	No	No	Yes	Yes

Notes: The table reports different measures of racial wage gaps based on employer-ascribed race restricting the sample to years before the candidate ever ran for an election. All columns include year fixed-effects. The first column always shows the unconditional relationship (apart from year fixed-effects), and the remaining columns progressively add controls. The final column includes all fixed-effects, as well as controls for self-declared race. Panel A shows coefficient estimates of OLS regressions of log hourly wage on employer-ascribed race. Panel B shows reduced form regressions of log hourly wage on skin tone. Panel C presents the first-stage coefficient from regressing employer-ascribed race on skin tone. Finally, Panel D shows 2SLS estimates of regressing log hourly wage on employer-ascribed race, instrumented by skin tone. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.14: Wage Disparities, Self-Declared Race

	(1)	(2)	(3)	(4)	(5)	(6)
[A] OLS: Conventional Disparity						
Non-White Self-Declared	-0.163***	-0.107***	-0.070***	-0.077***	-0.043***	-0.019***
	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
[B] Reduced Form: Colorism Disparity						
Skin Tone	-0.026***	-0.032***	-0.026***	-0.028***	-0.015***	-0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
[C] First Stage: Race Function						
Skin Tone	0.154***	0.149***	0.126***	0.125***	0.122***	0.115***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
[D] 2SLS: Perception-Normalized Disparity						
Non-White Self-Declared	-0.168***	-0.217***	-0.204***	-0.225***	-0.123***	-0.048***
	(0.008)	(0.008)	(0.009)	(0.009)	(0.008)	(0.008)
N	1,510,833	1,510,833	1,510,833	1,510,510	1,510,445	1,249,942
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	Yes	Yes	Yes	No
Industry FE	No	No	No	Yes	No	No
Industry × Occupation FE	No	No	No	No	Yes	No
Firm × Occupation FE	No	No	No	No	No	Yes

Notes: The table reports different measures of racial wage gaps based on self-declared race. All columns include year fixed-effects. The first column always shows the unconditional relationship (apart from year fixed-effects), and the remaining columns progressively add controls. The final column includes all fixed-effects, as well as controls for self-declared race. Panel A shows coefficient estimates of OLS regressions of log hourly wage on self-declared race. Panel B shows reduced form regressions of log hourly wage on skin tone. Panel C presents the first-stage coefficient from regressing self-declared race on skin tone. Finally, Panel D shows 2SLS estimates of regressing log hourly wage on self-declared race, instrumented by skin tone. Standard errors clustered at the individual level are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.15: Skin Tone Gaps for Repeated Candidates

	(1)	(2)	(3)	(4)	(5)	(6)
[A] Hourly Wage (log), OLS						
Skin Tone	-0.051***	-0.044***	-0.037***	-0.040***	-0.020***	-0.010***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
[B] Hourly Wage (log), Split-Sample-IV						
Skin Tone	-0.114***	-0.112***	-0.099***	-0.109***	-0.054***	-0.029***
	(0.007)	(0.007)	(0.008)	(0.007)	(0.006)	(0.008)
[C] Non-White Employer-Ascribed, OLS						
Skin Tone	0.152***	0.152***	0.119***	0.117***	0.113***	0.099***
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.003)
[D] Non-White Employer-Ascribed, Split-Sample-IV						
Skin Tone	0.346***	0.378***	0.315***	0.311***	0.307***	0.303***
	(0.005)	(0.007)	(0.006)	(0.006)	(0.006)	(0.014)
N	652,140	652,140	652,140	651,990	651,902	536,491
KP F-stat	6,651	5,147	4,290	4,224	4,020	959
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	Yes	Yes	Yes	No
Industry FE	No	No	No	Yes	No	No
Industry × Occupation FE	No	No	No	No	Yes	No
Firm × Occupation FE	No	No	No	No	No	Yes

Notes: The table reports results from regressions of log hourly wages (Panels A and B) and employer-ascribed race (Panels C and D) and on skin tone. The sample consists of all individuals who ran for political office multiple times and for whom multiple electoral photos are available. Panels A and C present OLS estimates, in which skin tone is measured by averaging across all photos of an individual. Panels B and D implement a split-sample IV approach, where the average skin tone from one randomly selected half of an individual's photos is instrumented with the average from the other half of photos. The first column only controls for year fixed-effects, and the remaining columns progressively add controls. Standard errors clustered at the individual level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.16: Characteristics of Compliers in Perception-Normalized Disparities Analysis

	(1) Complier mean	(2) Population mean
Worker Characteristics		
Male	0.758	0.747
	(0.004)	(0.001)
Age (years)	34.513	34.485
	(0.066)	(0.017)
Log Hourly Wage (2018 BRL)	2.138	2.176
	(0.005)	(0.001)
Highest Degree Attained		
Less than Elementary	0.025	0.027
	(0.001)	(0.000)
Elementary School	0.145	0.137
	(0.003)	(0.001)
Middle School	0.234	0.233
	(0.004)	(0.001)
High School	0.502	0.516
	(0.004)	(0.001)
Undergraduate or more	0.053	0.088
	(0.002)	(0.001)
Establishment Characteristics		
Log firm Size	4.239	4.095
	(0.019)	(0.004)
Average Log Hourly Wage	2.148	2.147
	(0.004)	(0.001)
Sector		
Agriculture	0.025	0.025
	(0.001)	(0.000)
Extractive industry	0.009	0.009
	(0.001)	(0.000)
Manufacturing	0.247	0.243
	(0.003)	(0.001)
Construction	0.052	0.063
	(0.002)	(0.000)
Retail	0.221	0.248
	(0.003)	(0.001)
Transport & communication	0.124	0.128
	(0.003)	(0.001)
Finance & real state	0.144	0.147
	(0.003)	(0.001)
Education & health	0.087	0.085
	(0.002)	(0.001)
Other services	0.051	0.051
	(0.002)	(0.000)
Region		
North	0.020	0.054
	(0.001)	(0.000)
Northeast	0.100	0.167
	(0.003)	(0.001)
Southeast	0.499	0.508
	(0.004)	(0.001)
South	0.110	0.194
	(0.003)	(0.001)
Center-West	0.047	0.077
	(0.002)	(0.001)

Notes: The table reports estimates of the average characteristics of compliers in the PND analysis of Table 3 (Panel D). To characterize compliers, we follow Hull (2025) and estimate the following IV model: $R^E_{it} \times \mathrm{X}_{it} = \alpha + \beta R^E_{it} + X_{it} + \lambda_t + \epsilon_{it}$, where R^E_{it} is the employer-ascribed race of worker i at time t, which we instrument with workers' skin tone, and X_{it} is the value of the characteristic of interest. The coefficient β identifies the mean of X_{it} among compliers, i.e., among individuals for whom employer-ascribed race changes with their skin tone. Each mean reported in Column (1) of the table is from one such IV regression and is compared with the mean across all individuals in our sample (Column 2). Standard errors are in parentheses.

C Details on Data

C.1 Employer-Ascribed Race

Information on employer-ascribed race, as well as many other employee and firm characteristics used in this paper, comes from the *Relação Anual de Informações Sociais* (RAIS - Annual Relation of Social Information). RAIS is an annual data collection instrument administered by Brazil's Ministry of Labor and Employment. Its objective is to "meet the needs of control, statistics, and information of government entities in the social and labor sectors" (Brasil, 2021). All formal firms (i.e., those registered with a tax identifier) and individuals who hire formal employees (excluding domestic employees) are required to report their employees' contract information to RAIS annually. Since 2003, one of the requested fields in RAIS is the employee's "race or color."

In RAIS, employers are responsible for reporting information on behalf of their employees. In larger firms, these records are typically submitted by a manager, HR department, or accountant, whereas in smaller firms the owner often completes the reporting. Figure C.1 shows the online interface, maintained by Brazil's Ministry of Labor and Employment, through which employers enter the information. Until 2018, RAIS reporting was conducted exclusively by the employer through this online system. The official RAIS completion manual from 2003 to 2018 provides no specific guidance on completing the race field, and there was no standardized training for those who entered the data.

Understanding how race is measured in RAIS thus requires examining how employers obtain information about workers' characteristics. At the time of hiring, workers must present various official documents, including the *Carteira de Trabalho e Previdência Social* (CTPS - Work and Social Security Card). The CTPS contains basic identifying information – including name, date of birth, gender, and residence – but does not include a field for race. Workers must also provide a photograph and proof of required education. Employers then record worker information in the *Livro de Registro dos Empregados* (LRE - Employee Registration Book), an internal employee record maintained at the establishment level. Employers must enter each new worker into the LRE, and the information recorded there is used to comply with various reporting requirements, including RAIS. The required information includes date of birth, hire date, identification number, and various job characteristics. There is also space for a photograph of the worker. Unlike the CTPS, the LRE commonly includes a field for race (cor, literally "color"), although completing this field is not mandatory (Cornwell et al., 2017).

Starting in 2019, the Brazilian government has integrated RAIS with other reporting forms that firms are required to complete within the E-Social platform. This changed how employee data is collected, since E-Social directly collects employees' information at the time of hire (which began to include mandatory self-declared race collection in 2023). Therefore, since 2019 and especially since 2023, employee race in RAIS is likely more closely aligned with their self-declared race at the time of hire. However, prior to 2019 (the period covered by our data), RAIS reporting is best conceptualized as employers' racial perceptions of their workers.

To verify this, we conducted qualitative interviews with employers and Human Resources (HR) managers with experience reporting RAIS data. We found that firms rarely asked employees directly for their self-identified race prior to 2019. We conducted eight interviews between October 2025 and March 2026: two with small-business owners employing fewer than 50 workers, and six with HR managers experienced in larger firms.[1] The two small-business owners reported that, until 2018, they filled out the RAIS form themselves and imputed employees' race based on their own perceptions, without asking employees to self-identify.

Practices varied more among larger firms. The two most common procedures described in the interviews, each mentioned roughly equally often, were to report race based on employee photos collected at hiring or to ask direct managers to provide employee information to the firm's central HR team. The HR managers we interviewed both reported using these practices in their own firms and described them as common in other firms during the period we studied. In a single case, a health-sector firm with approximately 10,000 employees, the interviewee reported that the firm collected self-declared race at hiring and used this information to complete the RAIS form. Even in this case, however, the interviewee explained that not all newly hired employees completed the form, and that HR staff would then complete the profile using employee photos.

Overall, the interviews suggest that race in RAIS during this period is better interpreted as reflecting employers' racial perceptions than employees' self-identification. This interpretation is consistent with the language used in technical reports issued by the Brazilian government. For instance, a 2003 report by the *Instituto de Pesquisa Econômica Aplicada* (IPEA, Institute for Applied Economic Research) states that, in RAIS, "it is usually not the employees who report their characteristics but the company" (Osório, 2003). Similarly, a more recent IPEA report states that, in RAIS, "the classification format is hetero-classification; that is, racial declaration is made by the employer or by an administrative employee responsible for it" (Silveira, 2022).

Figure C.1: RAIS Interface

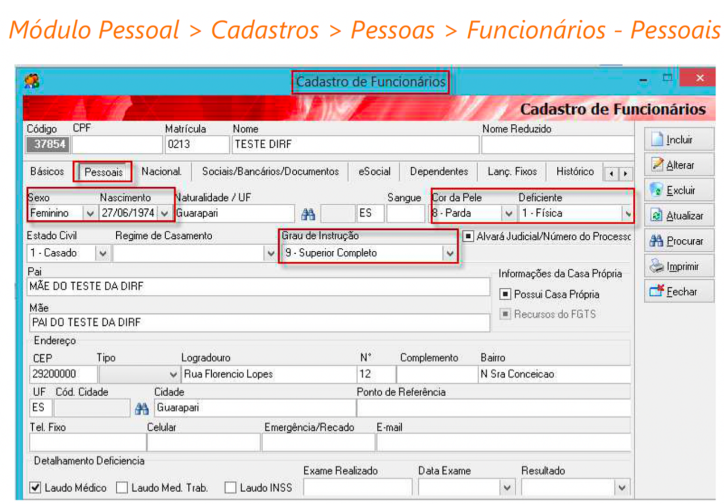

Notes: RAIS interface for reporting employee's information (2017 version).

[1] The two small firms are in the retail sector. The larger firms are in the health, education, retail, and finance sectors, and range from 300 to 20,000 employees.

C.2 Skin Tone

Electoral Photos. We assemble a dataset of nearly three million official photographs of candidates running for political office in Brazil between 2004 and 2024. These photographs are used during elections in Brazil's electronic voting machines, and they typically differ from campaign materials used in advertisements. Electronic voting was introduced by the *Tribunal Superior Eleitoral* (TSE) in 1996 with the aim to reduce the time and costs of vote counting. Voters cast their ballots by typing a candidate's identification number into the machine's keypad. Once the full number is entered, the interface displays the candidate's name, party label, and photograph, after which the voter may either confirm the vote or restart the process. The photograph is thus visible only at the moment of verification, used as a mechanisms to help low-literacy voters cast their ballot and reduce the incidence of misrecorded votes (Fujiwara, 2015).

The submission of a valid photograph is mandatory for a candidacy to be accepted by the TSE. Technical and stylistic guidelines for these photographs are specified in electoral resolutions and have evolved over time. Most recently, Resolution No. 23,609 (2019), Article 27 II, details the following requirements:

- Dimensions: 161 x 225 pixels (W x H); without frame,
- Color depth: 24bpp,
- Background: uniform color,
- Characteristics: frontal (bust) view, with attire suitable for official photography, ensuring the use of ethnic or religious clothing and body paint, as well as accessories necessary for the person with a disability; the use of scenic elements and other adornments is prohibited, especially those that have connotations of electoral propaganda or that induce or hinder the recognition of the candidate by the electorate.

Before that, Resolution No. 23,405 (2014), Article 27 III, imposed similar requirements but specified an 8bpp grayscale format with a uniform, preferably white, background.

Skin Tone Measurement. Next, we provide details on our skin tone measurement algorithm. We begin by detecting the skin part in each photo. For that, we use Google MediaPipe's Image Segmentation tool, which relies on a vision transformer neural network to locate different objects in an image and allows us to segment the skin from non-skin areas of the photo. Specifically, we apply the *multi-class selfie segmentation model*, which classifies pixels as *background*, *hair*, *clothing*, *accessories*, *face-skin*, and *body-skin*.[2] We combine the latter two categories to form the skin mask of each photo. Next, we identify the representative color of the skin area, closely following the algorithm of Adukia et al. (2023). In particular, we use k-means clustering to partition the pixel colors into distinct groups based on their RGB color values. We group the pixels into $k = 5$ clusters but only use the largest three clusters for the skin tone measure, as the smaller clusters often correspond to shadows and non-skin parts such as eyes or lips. For each retained cluster, we compute its centroid,

[2] https://ai.google.dev/edge/mediapipe/solutions/vision/image_segmenter.

which summarizes its characteristic color, and then convert it from the RGB color space to the $L^*a^*b^*$ color space.[3] To obtain a single representative skin color for each photo, we average the clusters' L^*, a^*, and b^* values, weighting each cluster by its assigned share of the skin-area pixels. In order to reduce the dimensionality of skin color, we focus on the darkness or lightness of the color – the "perceptual skin tint" (Adukia et al., 2023) – by retrieving only the L^* value from each photo's representative skin color.[4] Finally, we invert the L^* value, such that its scale ranges from 0 (lightest) and 100 (darkest). Figure C.2 shows examples of candidates with measured skin tones at the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentile of the distribution. For our analysis, we standardize the skin tone measure to have mean zero and unit standard deviation.

Figure C.2: Skin Tone Examples

Percentile	10	25	50	75	90
Skin tone	23.3	29.9	37.0	46.1	56.7
Race (self)	*Parda* Brown	*Branco* White	*Pardo* Brown	*Preta* Black	*Parda* Brown

Notes: The figure presents five candidate examples, showing their electoral photographs, the extracted skin tone with its corresponding position in the skin-tone distribution, and the candidates' self-reported race.

Comparison to Human-Assessed Skin Tone. As a validation exercise, we can compare our algorithmic skin tone measure to assessments made by a representative sample of human raters. For that, we draw on data collected by Bueno and Dunning (2017). They conducted an online survey of about 1,100 Brazilian respondents who were shown electoral photos of candidates running in the 2008 and 2010 elections and were asked to rate candidates' skin color on a scale from 1 (lightest) to 10 (darkest). We successfully match 4,826 candidates, each of whom was assessed by an average of three respondents. Appendix Figure C.3 presents a binned scatterplot for the relation between our algorithmic skin tone measure and the mean survey rating for the same photo, showing a strong positive relation

[3] The $L^*a^*b^*$ color space – also referred to as CIELAB – is a device-independent system defining colors with three values: L^* (lightness), a^* (red-green axis), and b^* (yellow-blue axis). Created by the International Commission on Illumination, the color space is perceptually uniform, meaning equal distances between colors represent similar perceived differences.

[4] In the figures in the paper that depict color differences across the skin tone distribution, we always vary the L^* value for a skin tone with $a^* = 17.2$ and $b^* = 19.1$, which correspond to the sample averages of a^* and b^* across all our colored photos.

across the distribution. The overall correlation between both measures is 0.39. This can be compared to how much human raters diverge in their evaluation of the same photo. The average inter-rater correlation is 0.37. Hence, our algorithm's agreement with the average human perception is on par with the level of agreement observed among humans themselves.

Figure C.3: Algorithmic vs. Survey-assessed Skin Tone

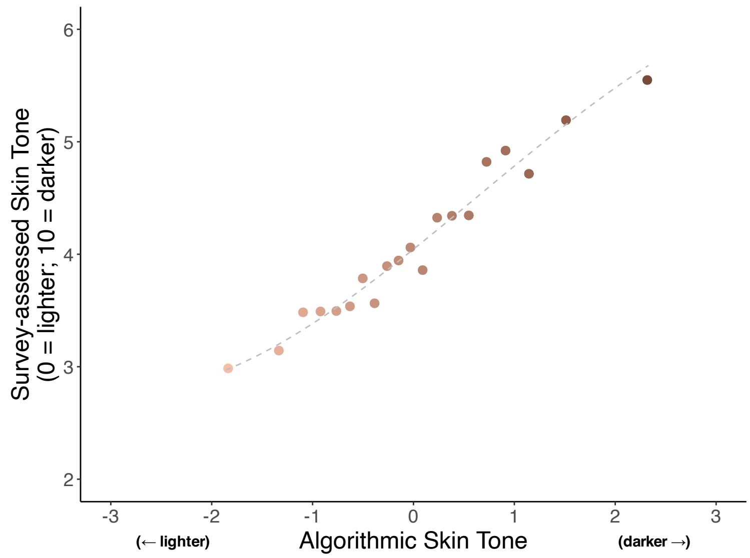

Notes: The figure shows a binned scatterplot of survey-assessed skin tone and our algorithmic skin tone measure. The former is obtained from Bueno and Dunning (2017) and captures the average rating across Brazilian human coders who assess the skin tone of a given political candidate on a scale from 1 (lightest) to 10 (darkest). $N = 4,826$ candidates. The plot is constructed following the procedure in Cattaneo et al. (2024). For each sample, we divide the algorithmic skin tone distribution in twenty equal-sized bins and, for each bin, plot the average of survey-assessed skin tone.

Colored vs. Gray-scale Photos. About half of all candidate photos in our estimation sample (47.1%) are in color, with the other half – especially photos from candidacies in earlier years – being in grayscale.[5] Note that the algorithm is suitable for extracting skin tint from both types of images (following Adukia et al. (2023)). To assess potential differences in measurement quality, in Table C.1 we examine correlations with employer-ascribed race separately for skin tone measured from colored and grayscale photos.[6] In column (1), we replicate our main result from Section 5 when using all available photos of each candidate, while in columns (2) and (3), we only use colored or grayscale photos to measure the skin tone of each candidate (the corresponding binned scatterplots are shown in Figure C.4). For 63.2% (62.2%) of all candidates, we have at least one colored (grayscale) photo (remember that we average the skin tone across photos for individuals who run multiple times), which

[5]In particular, all photos of candidates until 2016 were in grayscale, in 2018 and 2020 they started to be published in color (depending on location and party), and since then almost all photos are colored.

[6]Note that Bueno and Dunning (2017) have examined photos of candidates running in 2008 and 2010, which are all in grayscale, such that we cannot analyze potential differences in correlations with human-assessed skin tone.

explains the reduction in sample size in the two columns. We find that the predictive power of skin tone for racial ascription is only modestly higher when using colored photos rather than grayscale photos, which can be explained by larger confounding influence from lighting in grayscale photos. At the same time, comparing columns (1) and (2), we see that the predictive power of colored photos (column (2)) is almost the same as when using all photos. This is because aggregating over more photos to measure the skin tone of each individual improves the measurement, which is highlighted in columns (4) to (6). Here, we focus on candidates who run multiple times and for whom we observe at least one colored and one grayscale photo. As can be seen, averaging across colored and grayscale photos yields a substantially larger correlation with employer-ascribed race compared to using colored or grayscale photos alone. Overall, Table C.1 shows that incorporating both colored and grayscale photos in our skin tone measurement increases the sample of candidates that we can analyze without compromising measurement quality. Moreover, note that our estimation of perception-normalized disparities (PNDs) in Section 6 precisely accounts for the predictive power of skin tone for racial ascription, adjusting skin tone disparities for any potential noise in how skin tone is measured and how it translates into racial ascriptions.

Table C.1: Relationship between Employer-Ascribed Race and Skin Tone, Colored vs. Grayscale Photos

	All Candidates			Candidates with Colored and Grayscale Photos		
	(1)	(2)	(3)	(4)	(5)	(6)
Skin Tone - All Photos	0.107***			0.164***		
	(0.001)			(0.002)		
Skin Tone - Colored Photos		0.111***			0.114***	
		(0.001)			(0.002)	
Skin Tone - Grayscale Photos			0.090***			0.094***
			(0.001)			(0.002)
N	1,480,624	939,770	922,032	381,178	381,178	381,178

Notes: The table reports results from regressions of employer-ascribed race (indicator equal to one if the worker is ascribed non-White) on skin tone measured from all candidate photos vs. only colored photos vs. only grayscale photos. All skin tone measures are standardized using the mean and standard deviation from the distribution of skin tone based on all photos. Columns (1) to (3) use the full sample of candidates, while columns (4) to (6) restrict the sample to candidates who run for office multiple times and for whom we observe at least one colored and at least one grayscale photo. All regressions include year fixed effects. Standard errors clustered at the individual level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure C.4: Relationship between Employer-Ascribed Race and Skin Tone, Colored and Grayscale Photos

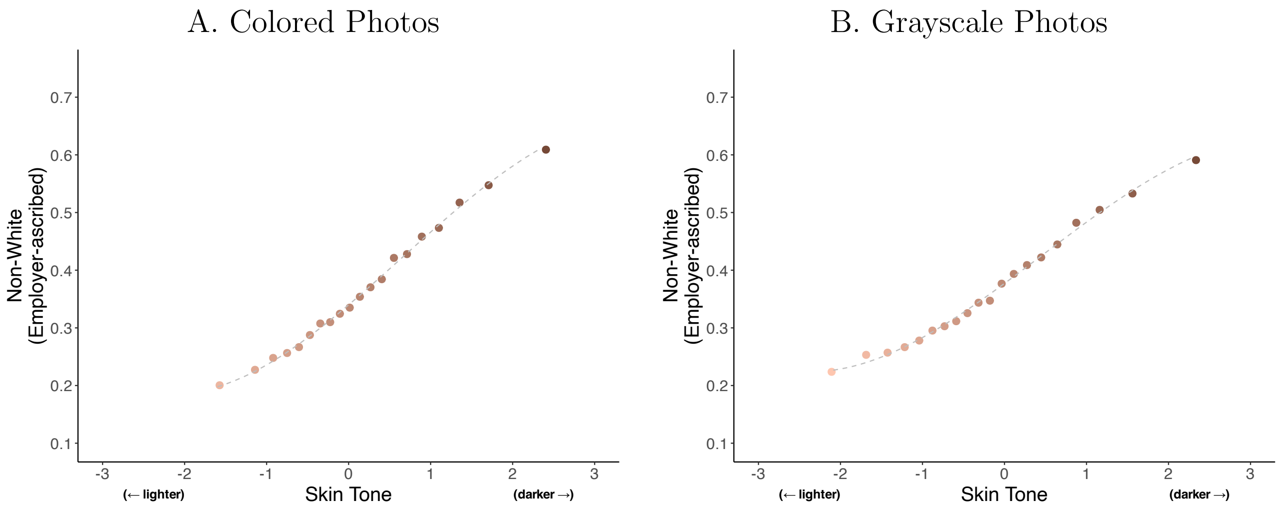

Notes: The figures show binned scatterplots of employer-ascribed race (an indicator equal to one if the worker is ascribed non-White) and skin tone measured from colored photos vs. grayscale photos (corresponding to columns (2) and (3) of Table C.1 but without year fixed effects). All skin tone measures are standardized using the mean and standard deviation from the distribution of skin tone based on all photos. The plots are constructed following the procedure in Cattaneo et al. (2024). For each sample, we divide the skin tone distribution in twenty equal-sized bins and, for each bin, plot the average of employer-ascribed race. We fit third-degree polynomials based on the underlying data of each sample.

C.3 Machine-Learning Prediction of Visual Cues

In addition to skin tone, we construct a flexible summary of the racial cues visible in the candidates' photos. For that, we use a machine-learning algorithm to predict employers' racial ascription based on the photos. Following recent computer vision papers on ethnicity inference from face images (Soudbakhsh et al., 2026; Fournier-Montgieux et al., 2025a), we apply a transfer learning approach that integrates a pretrained face recognition encoder with a non-linear classification head that we train on our data. Several studies have documented benefits of transfer learning in face analysis tasks relative to full end-to-end model training (Fournier-Montgieux et al., 2025b; Narayan et al., 2025; Sun et al., 2024; Amato et al., 2019).

For the pretrained face recognition model, we use the *Facenet-512d* encoder (Schroff et al., 2015), which is the best-performing encoder in face recognition benchmarks among those wrapped in the *DeepFace* Python package developed by Serengil and Ozpinar (2024). The encoder extracts a 512-dimensional vector of facial embeddings from each candidate photo. Before applying the face recognition encoder, we segment faces from the photos, relying on the *OpenCV* face detector that is also available in the *DeepFace* package.

We then use the facial embeddings, together with our skin tone measure, to train an RBF-kernel support vector machine (SVM) that predicts whether employers ascribe an individual as non-White.[7] Specifically, the outcome is the mode of employers' racial ascription across all years in which the individual appears in RAIS. We estimate the SVM in R using the *kernlab* implementation through the *tidymodels* interface. We use a random training dataset of 50,000 photos, with the remaining 510,849 photos used as hold-out data. Hyperparameters are tuned with a grid search, using 50% of the training sample to fit candidate models and using 50% as validation sample. After selecting the parameter combination with the highest validation ROC-AUC, we refit the SVM on all 50,000 training photos. The selected parameters are RBF kernel sigma = 0.0001 and regularization cost = 10. The model yields predicted probabilities of being ascribed non-White by the employer. In the hold-out data, we obtain a ROC-AUC of 0.736, indicating that the model has a 73.6% probability to correctly to assign a higher prediction to a worker ascribed as non-White by the employer than to a worker ascribed as White. The R^2 of the predictions in the hold-out data is 0.147. We only use the predictions for hold-out photos in the estimation in Tables B.10 and B.11.[8]

[7] We also experimented with a random forest model, which had notably worse prediction performance.

[8] For individuals with multiple photos, i.e. those running for office multiple times, we average the predicted probabilities across photos.

D AKM Model

D.1 Empirical Strategy

Estimation Model. This section describes how we leverage workers' moves across establishments to quantify the extent to which ascribed race in RAIS varies depending on the employer, in an exercise in the spirit of Abowd et al. (1999). We estimate the following model using the full RAIS sample:

$$y_{it} = \alpha_i + \psi_{\boldsymbol{j}(i,t)} + X_{it}'\beta + \varepsilon_{it} \tag{1}$$

where y_{it} is an indicator equal to one if worker i at year t is ascribed as non-White (*preto* or *pardo*), and $\boldsymbol{j}(i,t) \in \{1, \ldots, J\}$ is an identity function that returns the firm j at which worker i was employed during year t.

The fixed effects model above divides a worker's employer-ascribed race into four components. First, α_i, the person's race effect, represents the portion of a worker's racial identity that all employers interpret equally. It captures features unambiguously associated with a specific racial group, such as certain skin tones, as well as signals that all employers associate with a particular racial group. Second, the firm's race effect $\psi_{\boldsymbol{j}(i,t)}$ represents the portion of ascribed race that is specific to each employer. This effect may include differences in employers' diligence in race reporting, differences in how employers interpret racial cues, and differences in how workers tend to present to a particular employer. Third, X_{it} is a set of the following time-varying controls: age squared, age cubed, and year, all interacted with workers' education (5 educational levels).[9] Finally, the error term ε_{it} captures all other match-specific factors affecting ascribed race, including idiosyncratic measurement error.

Variance Decomposition. Under the strict exogeneity assumption (discussed below), the person and firm effects in Equation (1) are identified. Estimating the above model allows us to decompose the variance of employer-ascribed race in RAIS into a component that is specific to firms, another that is specific to workers, and their correlation, as proposed by Abowd et al. (1999) in the case of wages:

$$\begin{aligned}\mathbb{V}[y_{it}] = {} & \mathbb{V}[\alpha_i] + \mathbb{V}\left[\psi_{\boldsymbol{j}(i,t)}\right] + \mathbb{V}[X_{it}'\beta] + \mathbb{V}[\varepsilon_{it}] \\ & + 2\mathbb{C}\left[\alpha_i, \psi_{\boldsymbol{j}(i,t)}\right] + 2\mathbb{C}[\alpha_i, X_{it}'\beta] + 2\mathbb{C}\left[\psi_{\boldsymbol{j}(i,t)}, X_{it}'\beta\right]\end{aligned}$$

where $\mathbb{V}$ and $\mathbb{C}$ are the finite-sample variance and covariance operators, respectively. Following Kline (2024), we present our decomposition in terms of the covariance-adjusted variance in ascribed race, i.e., we scale the variance and covariance of α_i and $\psi_{\boldsymbol{j}(i,t)}$ by $\mathbb{V}[y_{it} - X_{it}'\beta]$.

[9] To separately identify the person-effects and the time-varying controls, it is necessary to make a normalization assumption. Following Gerard et al. (2021) and Card et al. (2018), we assume that $X_{it}'\beta = 0$ for 40-year-olds. This implies that we measure person effect α_i for 40-year-old workers, which is approximately the age where the average worker achieves the peak of their wage in our sample. Differently from the literature estimating mover's design models outside of the labor context (e.g., Cantoni and Pons, 2022; Finkelstein et al., 2016), we do not control for relative move time because it is too computationally intensive, there is no strong evidence of pre-trends without controlling for them, and this is the usual practice in the labor literature.

Identification Assumption and Specification Tests. OLS estimates of Equation (1) will yield unbiased estimates of the person and establishment fixed effects under the assumption of strict exogeneity of this error term, i.e., if $\mathbb{E}\left[\varepsilon_{it}|\boldsymbol{j}(i,s)=j, X_{is}=x\right]=0$ for all workers $i \in \{1,\ldots,N\}$ and all periods $(s,t) \in \{1,\ldots,T\}^2$.

The strict exogeneity identification assumption has two key implications in our context (which mirror the traditional AKM analysis as discussed, for instance, in Kline, 2024). First, worker mobility across firms must be uncorrelated with the time-varying residual component of ascribed race. This assumption is often referred to as *exogenous mobility.* In our setting, it rules out that changes in underlying racial identities are correlated with moves to particular types of firms. For instance, workers who become more likely to be ascribed non-White cannot become more likely to move to *darkening* firms.

As in previous work using similar models to study variation in wages (Gerard et al., 2021; Card et al., 2013), we test this hypothesis by performing an event-study analysis of changes in the ascribed race of workers moving across different types of firms. We split the sample of firms into quartiles of the estimated firm race effects $\hat{\psi}_{\boldsymbol{j}(i,t)}$, and consider the changes in ascribed race for workers moving from firms in the top and bottom quartile to firms in any quartile of the distribution. For that, we follow previous literature and restrict the sample to workers who stayed in the origin job for at least 2 years before the move. Results are displayed in Figure D.1. Reassuringly, we find no evidence of differential trends in ascribed race before the move, supporting the assumption that mobility is exogenous to time-varying components of the residual race change (however, we cannot rule out that sudden changes in racial identity are causing moves.)

Figure D.1: Event Studies around Job Moves

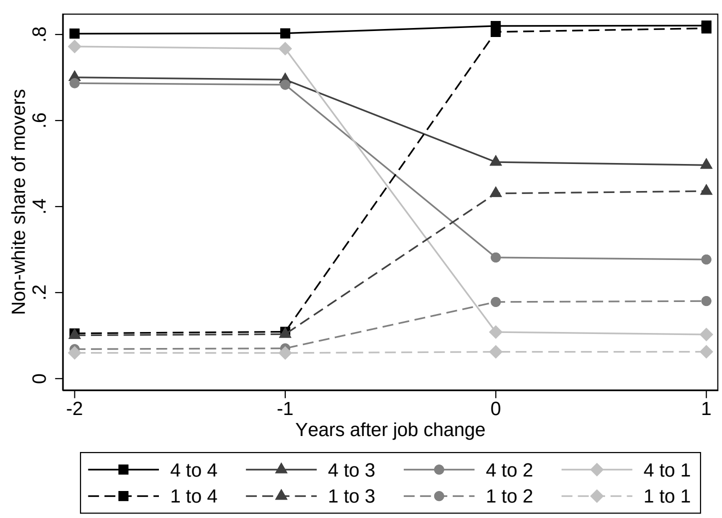

Notes: The figure displays the likelihood of being reported non-White by the employer among workers who moved from establishments in the top and bottom quartiles of establishment race effects to destination establishments in any of the other quartiles. Movers are defined as workers who separated from their origin establishment from 2003-2013, were reemployed in the destination establishment the following year, and were employed in both the origin and destination establishment for at least 2 years.

The event study plot of Figure D.1 also provides helpful evidence on the second implication of strict exogeneity: that ascribed race is additively separable among firm and worker effects. Additive separability implies that effects of moves from *whitening* to *darkening* firms and vice-versa must be symmetric, and this is exactly what we observe: workers moving from firms in the top quartile (i.e., firms that tend to *darken* workers) experience a change in the likelihood of being ascribed non-White that is similar in magnitude and opposite in sign to workers moving from firms from the bottom quartile (i.e., firms that tend to *whiten* workers).

D.2 Results

We estimate model (1) in the largest connected set of workers. As in any mover design analysis, worker and firm effects are only identified within the set of establishments linked by worker moves. We first focus on the largest connected set of firms in our sample, which comprises 99.1% of worker-year observations in the full sample. Results are shown in the first column of Table D.1. The model fits well, with an adjusted R^2 of 76.4% and a root mean squared error (RMSE) of 0.234.

We find that firm-specific reporting behavior accounts for a substantial share of the variance in ascribed race observed in RAIS. The standard deviation of our outcome (an indicator equal to one for workers being ascribed non-White) is 0.482. In comparison, the standard deviation of the person effects is 0.332, while the standard deviation of establishment effects is 0.263. The implied variance decomposition shows that person effects explain about 47.2% of the variance in race in RAIS. Variation in establishment-specific race reporting explains 29.6% of the variance.[10] The remaining variance is largely residual variation, as the covariance between person and firm effects is low. Figure D.2 displays the distribution of the estimated firm and person race effects.[11]

Limited mobility of workers may lead to biased estimates of the variance and covariance of firm and person effects. Thus, we also report bias-corrected estimates in the leave-one-out-connected set of firms and workers, following the approach of (Kline et al., 2020) (KSS). As shown in the second column of Table D.1, once we account for limited mobility using the KSS decomposition, this results in even smaller estimates for the share of variance explained by person effects (30.8%), while the share explained by firm effects changes only mildly (25.5%). Moreover, the correlation between person and firm effects gains in importance, now explaining 9.1% of the variation. Overall, these results suggest an even larger role for establishments in shaping racial ascriptions (when jointly considering firm effects and their covariance with person effects).

These findings are important and novel. First, the worker-specific, portable component of race explains at most half of the variation in ascribed race, indicating that race is not an innate characteristic of workers that is equally perceived across employers. Second, about

[10] Following Kline (2024), we report the decomposition in terms of the covariate-adjusted variance in employer-ascribed race, though results are remarkably similar if we reported unadjusted shares (since the standard deviation of the variation explained by our covariates is low).

[11] As an alternative exercise, we estimated the AKM model separately by gender. Correlating the firm effects estimated in the female and male sample, we find a coefficient of 0.55. The correlation suggests that firms' race reporting applies similarly to different types of workers, providing additional evidence in favor of the additive separability assumption discussed above.

one-third of the variation in ascribed race is explained by establishment variation, highlighting that firms vary widely in their race perceptions. For comparison, we use the same sample and estimate an AKM model for log hourly wages, with results shown in Table D.2. Similar to the results in (Gerard et al., 2021), firm effects account for 16.7% of wage variation. Thus, firms appear to be even more important in explaining differences in racial ascription than in wage setting.

Figure D.2: Distribution of Estimated Person and Establishment Race Effects

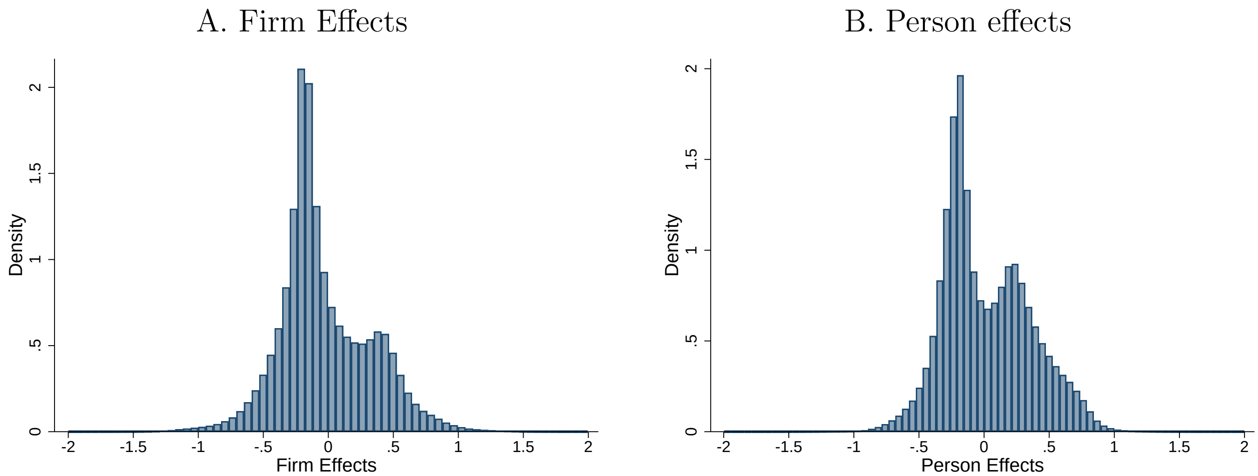

Notes: The figures display histograms of the estimated person and establishment effects for the two-way fixed effect models of employer-ascribed race.

Table D.1: Summary of Estimated Two-Way Fixed Effects Models

	Connected Set	
	Largest (1)	Leave-One-Out (2)
Percentage of sample (across person-year obs.)	99.1	95.5
Percentage non-White (across person-year obs.)	36.8	37.0
Std. deviation of non-White (across person-year obs.)	0.482	0.483
AKM Decomposition		
SD of person effects (across person-year obs.)	0.332	0.324
SD of establishment effects (across person-year obs.)	0.263	0.257
Correlation of person and establishment effects	0.070	0.069
Adjusted R^2 of the model	0.764	0.815
Root Mean Squared Error of the model	0.234	0.200
Percentage of covariate-adjusted variance in employer-ascribed race due to:		
Person effect	47.2	45.0
Establishment effect	29.6	28.3
Covariance of person and establishment effects	4.1	4.9
KSS (Bias Corrected) Decomposition		
SD of person effects (across person-year obs.)	-	0.268
SD of establishment effects (across person-year obs.)	-	0.244
Correlation of person and establishment effects	-	0.162
Percentage of covariate-adjusted variance in employer-ascribed race due to:		
Person effect	-	30.8
Establishment effect	-	25.5
Covariance of person and establishment effects	-	9.1
Number of establishments	4,882,444	3,562,840
Number of workers	49,289,418	47,496,834
Number of person-year observations	294,255,726	284,420,494
Percentage of workers with any job change	66.8	66.6
Percentage of workers with any race change	29.5	29.5

Notes: The table summarizes the results from estimating the two-way fixed effect model (1), where the outcome is an indicator equal to one if worker i at year t is ascribed as non-White (*preto* or *pardo*) by their main employer. The models include worker and establishment fixed effects, alongside year dummies interacted with five educational level dummies, and quadratic and cubic age interacted with the educational level dummies. Column (1) summarizes the model estimated on the largest connected set of firms in our RAIS sample from 2003-2015. Column (2) summarizes the model estimated on the largest leave-one-out connected set of moves. For this model, we show results both of the simple AKM model, and bias-corrected estimates following Kline et al. (2020). The standard deviation and correlation estimates of the KSS Decomposition are already covariate-adjusted.

Table D.2: Summary of Estimated Two-Way Fixed Effects Models for Log Hourly Wages

	Connected Set
	Largest (1)
Percentage of sample (across person-year obs.)	99.1
Percentage non-White (across person-year obs.)	36.8
Std. deviation of log-hourly wage (across person-year obs.)	0.679
AKM Decomposition	
SD of person effects (across person-year obs.)	0.514
SD of establishment effects (across person-year obs.)	0.285
Correlation of person and establishment effects	0.371
Adjusted R^2 of the model	0.913
Root Mean Squared Error of the model	0.200
Percentage of covariate-adjusted variance in wages due to:	
Person effect	54.2
Establishment effect	16.7
Covariance of person and establishment effects	22.4
Number of establishments	4,882,444
Number of workers	49,289,418
Number of person-year observations	294,255,726
Percentage of workers with any job change	66.8
Percentage of workers with any race change	29.5

Notes: The table summarizes the results from estimating the two-way fixed effect model (1), where the outcome is the log hourly wages of worker i at year t. The models include worker and establishment fixed effects, alongside year dummies interacted with five educational level dummies, and quadratic and cubic age interacted with the educational level dummies. The model is estimated on the largest connected set of firms in our RAIS sample from 2003-2015.

Appendix References

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Adukia, Anjali, Alex Eble, Emileigh Harrison, Hakizumwami Birali Runesha, and Teodora Szasz, "What We Teach About Race and Gender: Representation in Images and Text of Children's Books," *The Quarterly Journal of Economics*, 2023, *138* (4), 2225–2285.

Amato, Giuseppe, Fabrizio Falchi, Claudio Gennaro, Fabio Valerio Massoli, Nikolaos Passalis, Anastasios Tefas, Alessandro Trivilini, and Claudio Vairo, "Face verification and recognition for digital forensics and information security," in "2019 7th International Symposium on Digital Forensics and Security" IEEE 2019, pp. 1–6.

Brasil, "Decreto nº 10.854, de 10 de novembro de 2021," November 10 2021. Regulamenta disposições relativas à legislação trabalhista e institui o Programa Permanente de Consolidação, Simplificação e Desburocratização de Normas Trabalhistas Infralegais e o Prêmio Nacional Trabalhista, alterando o Decreto nº 9.580/2018. Diário Oficial da União, Seção 1, 11 nov. 2021.

Bueno, Natália S and Thad Dunning, "Race, Resources, and Representation: Evidence from Brazilian Politicians," *World Politics*, 2017, *69* (2), 327–365.

Cantoni, Enrico and Vincent Pons, "Does context outweigh individual characteristics in driving voting behavior? Evidence from relocations within the United States," *American Economic Review*, 2022, *112* (4), 1226–1272.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

_ , Jörg Heining, and Patrick Kline, "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 2013, *128* (3), 967–1015.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," *American Economic Review*, 2024, *114* (5), 1488–1514.

Cornwell, Christopher, Jason Rivera, and Ian M. Schmutte, "Wage Discrimination When Identity Is Subjective: Evidence from Changes in Employer-Reported Race," *Journal of Human Resources*, 2017, *52* (3), 719–755.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of geographic variation in health care: Evidence from patient migration," *The Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.

Fournier-Montgieux, Alexandre, Hervé Le Borgne, Adrian Popescu, and Bertrand Luvison, "Reliable and Reproducible Demographic Inference for Fairness in Face Analysis," *arxiv preprint arXiv:2510.20482*, 2025.

—, Michael Soumm, Adrian Popescu, Bertrand Luvison, and Hervé Le Borgne, "Fairer analysis and demographically balanced face generation for fairer face verification," in "2025 IEEE/CVF Winter Conference on Applications of Computer Vision" IEEE 2025, pp. 2788–2798.

Fujiwara, Thomas, "Voting Technology, Political Responsiveness, and Infant Health: Evidence From Brazil," *Econometrica*, 2015, *83* (2), 423–464.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative matching or exclusionary hiring? The impact of employment and pay policies on racial wage differences in Brazil," *American Economic Review*, 2021, *111* (10), 3418–3457.

Hull, Peter, "One Weird Trick to Characterize Effective Populations in Design-Based Specifications," Technical Report, Metrics Note 2025.

Kline, Patrick, "Firm wage effects," in Christian Dustmann and Thomas Lemieux, eds., *Handbook of Labor Economics Vol. 5*, 2024, pp. 115–181.

—, Raffaele Saggio, and Mikkel Sølvsten, "Leave-out estimation of variance components," *Econometrica*, 2020, *88* (5), 1859–1898.

Narayan, Kartik, Vibashan VS, and Vishal M. Patel, "FaceXBench: Evaluating multimodal LLMs on face understanding," *arxiv preprint arXiv:2501.10360*, 2025.

Osório, Rafael Guerreiro, "O sistema classificatório de cor ou raça do IBGE," Technical Report, Instituto de Pesquisa Econômica Aplicada (IPEA) 2003.

Schroff, Florian, Dmitry Kalenichenko, and James Philbin, "FaceNet: A unified embedding for face recognition and clustering," in "Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition" 2015, pp. 815–823.

Serengil, Sefik and Alper Ozpinar, "A Benchmark of Facial Recognition Pipelines and Co-Usability Performances of Modules," *Journal of Information Technologies*, 2024, *17* (2), 95–107.

Silveira, Leonardo, "Imputação da informação de raça/cor na Rais para o setor público brasileiro," Technical Report, Instituto de Pesquisa Econômica Aplicada (IPEA) 2022.

Soudbakhsh, Hasti, Sonjoy Ranjan Das, Bilal Hassan, and Muhammad Farooq Wasiq, "Transfer Learning-Based Ethnicity Recognition Using Arbitrary Images Captured Through Diverse Imaging Sensors," *Sensors*, 2026, *26* (3), 886.

Sun, Haomiao, Mingjie He, Tianheng Lian, Hu Han, and Shiguang Shan, "Face-MLLM: A large face perception model," *arxiv preprint arXiv:2410.20717*, 2024.