Racial Disparities in Labor Outcomes: The Effects of Hiring Discrimination over the Business Cycle *

Florian Kuhn† Luis Chanci‡

November, 2023

Abstract

The resume audit literature provides strong evidence of discriminatory practices in hiring, raising critical concerns regarding equitable labor market outcomes. While the impacts of these practices on disparities in labor market levels are better understood, their cyclical effects are less known. In this paper, we research how hiring discrimination affects the volatility of labor market outcomes for disadvantaged groups by integrating empirical findings from audit studies into a search-and-matching model with a modified urn-ball matching function. Intuitively, in recessions, there are more applicants per job opening, which hurts discriminated groups. Applying this model to the U.S. economy, we find that it accounts for approximately 70% of the excess business cycle volatility in the unemployment rates of African Americans, as recorded in CPS data. Our research highlights the broader economic implications of discrimination, stressing the necessity for policy interventions, and offers a novel framework for future studies on labor market inequalities.

JEL codes: E24, E32, J64, J71.

Keywords: Unemployment, Discrimination, Business Cycle.

---

*This paper was previously circulated under the title 'The Effects of Hiring Discrimination over the Business Cycle.' We extend our gratitude to the two referees and the co-editor of this journal, Prof. Sushanta Mallick, for their insightful comments that significantly enhanced the content and presentation of this paper. Additionally, our research on this project has greatly benefited from the valuable feedback of Christopher Hanes, Subal Kumbhakar, and participants at seminars held at Binghamton University, the Society for Economic Dynamics meeting, and the Encuentro de la Sociedad de Economía Chilena (SECHI).

†Department of Economics, Binghamton University. Email: fkuhn@binghamton.edu

‡Escuela de Ingeniería Comercial, Facultad de Economía y Negocios, Universidad Santo Tomás, Chile. Email: luischanci@santotomas.cl.

1 Introduction

Discrimination in hiring practices is not just a moral dilemma but also a critical concern in economics, particularly for its impact on vulnerable demographic groups and the underutilization of talent and skills. Although the literature of resume-based audit studies has highlighted racial disparities in the hiring process, the broader macroeconomic implications of these discriminatory practices, especially when bad economic times prevail, still remain underexplored. In this paper, we examine how the effects of hiring discrimination on labor market outcomes vary over the business cycle.

We start from the observation that correspondence or resume studies, due to their quasi-experimental nature, provide strong evidence for discrimination against certain demographic or social groups during the early stages of the hiring process. Many resume studies have shown that members of such groups face lower callback rates when applying for job openings. By design, these studies vary only the group status of a fictitious applicant and hold all other characteristics constant. The idea is that in this way, it is possible to identify a direct effect of group status on a labor market outcome (in these settings, often callbacks for interviews) as opposed to picking up an indirect effect of a variable that may be correlated with group status like, for example, education.

In our research design, we take the level of hiring discrimination as given. Staying as close as possible to the evidence from the quasi-experimental studies, as a measure of discrimination, we take the difference in conditional hiring rates; that is, the relative likelihood of getting hired from the same applicant pool for two otherwise identical workers. Thus, to investigate the cyclical implications of such different hiring rates, we use a search-and-matching model, the workhorse model for labor markets in macroeconomics, and embed a modified urn-ball matching function that allows for arbitrary degrees of hiring discrimination. Over the business cycle, the model predicts that the discriminated group suffers from higher unemployment volatility; when the economy enters a recession, the unemployment rate among discriminated workers increases more strongly. The intuition of the mechanism is that in recessions there are many candidates for each job opening, resulting in increased competition between workers of different groups, which in turn hurts the discriminated workers.

We apply the model to the U.S. context, as it is one of the most racially diverse countries, has been at the forefront of implementing policies aimed at reducing racial disparities, and has seen a resurgence of social movements in recent times.[1] Thus, in addition to the model, we employ the Current Population Survey (CPS) data to empirically examine the volatility of unemployment and job-finding rates for two particular demographic groups, women and blacks, since many resume studies investigate the degree of hiring discrimination for those groups.

While the statistics we present using CPS may not be entirely novel in the literature, they serve a twofold crucial purpose. Firstly, to see - empirically - if these volatilities are larger

[1] To illustrate, the Black Lives Matter movement brought the discussion about systemic racial inequities to the forefront.

than their counterparts for the groups of whites and males, respectively, and if so, by how much. Importantly, here we focus on the conditional employment and job-finding rates; that is, we calculate the respective probability of being employed and finding a job controlling for many observable characteristics. Thus, in line with previous literature, we encounter substantial differences in unemployment volatility for blacks compared to whites; conversely, we find no evidence of extra volatility for women relative to men. These findings are consistent with the fact that resume studies find a much stronger degree of discrimination for blacks than for women in the hiring process. Secondly, these empirical findings bolster the assessment of our theoretical model's performance.

We calibrate the search-and-matching model with two types of workers to quantitatively study the equilibrium effects of different hiring rates. The way we set up the matching function allows us to use the difference in conditional hiring discrimination as a parameter. In our baseline calibration, we take this value directly from the point estimates for the differential in callback rates found in resume studies. More generally, the model provides a mapping between the degree of discrimination and the volatility of labor market outcomes so that we can assess the order of magnitude of the effect that hiring discrimination has on cyclical labor market outcomes. In the calibration, which matches differences in mean employment and job-finding rates, the model indicates that discrimination rates at a level found in resume studies could account for over half of the heightened volatility in unemployment rates among Blacks throughout the business cycle.

Throughout our theoretical analysis of the business cycle effects, we hold fixed the intensity of hiring discrimination prevailing in the market. By this, we mean that the likelihood for an employer to hire a member of a disadvantaged group remains constant over time conditional on the size and makeup of the applicant pool. For example, consider the case where there are only two applicants to a job opening, a white and a black applicant, who are equal in all other characteristics observable to the employer. We will define as the degree of discrimination the relative likelihood of the two applicants to receive the job -in other words, how much likelier is it that the white applicant gets hired compared to the black applicant? This measure corresponds directly to the object of interest in resume studies, where the goal usually is to estimate a relative likelihood of receiving an interview callback. The most direct interpretation for why employers discriminate in the model is on the basis of taste, but as we again discuss in more detail below, we have a strong conjecture that the mechanism also works in an environment where hiring bias results from statistical discrimination.

An obvious direct implication of a lower likelihood of getting hired is a level effect leading to worse average outcomes for the disadvantaged group like a higher mean unemployment rate and longer expected unemployment spells. But in addition, and in the focus of this paper, there are dynamic effects over the business cycle: In recessions, the unemployment rate among black workers increases stronger than the unemployment rate for whites. The basic intuition of this is as follows. The labor market is slack in a recession, with many applicants per job opening. This increased competition for jobs is particularly bad for the discriminated group: Under the mechanism considered here, hiring discrimination has an effect whenever workers of the two groups compete for the same job. Since recessions are times of larger applicant pools, the odds are high that a majority worker will be picked over a minority worker as a result of discrimination. The result is a bigger drop in employment and in the job-finding rate for the disadvantaged group during recessions. We model this effect formally by extending Blanchard and Diamond (1994)'s urn-ball model for flexible rates of discrimination, nesting their model as a special case and taking it to a dynamic setting.

We remark that resume audit studies literature identifies discrimination in callback rates, not in hiring, which would be an unobservable outcome. Thus, our assumption that the discrimination documented in audit studies shows up in hiring is aimed at considering callbacks as a reflection of discrimination in a very early stage of the hiring process. In other words, any potential difference between the two rates does not take away the fact that one group is negatively affected in the early stages of the hiring process by not even being contacted or called to participate. Hence, our goal is to go further in the discussion and question the potential effects of discrimination over the business cycle. While some effects of hiring discrimination on an individual's labor market outcomes have been well studied both empirically and theoretically, the impact on macro outcomes such as unemployment volatility has been less investigated in the literature.

Also, our focus on the hiring margin does not rule out the existence of other types of discrimination, for example on the margin of wages, job separations, promotions, etc. In fact there is a broad literature assessing the importance and the consequences of many of these alternative channels (e.g., see the reviews by Lang and Lehmann, 2012; Fang and Moro, 2011). In this paper we mainly isolate the hiring channel because we have a relatively clear idea of its order of magnitude from the reduced-form evidence of resume studies, which allows us to study its effects quantitatively in the model – in particular its effect on unemployment volatility.

In the empirical part, we use the CPS to study the labor market outcomes for two of the main groups that resume studies have focused on: blacks and women. While not the only groups for which such resume studies have been conducted[2], they are the most straightforward to work with empirically in terms of data availability, definition of group membership, and exogeneity of group membership. In the empirical analysis we control for many individual characteristics that are also observable by employers. We find that unemployment rates exhibit excess volatility over the business cycle for blacks compared to whites. For example, given a five percentage point increase in the aggregate unemployment rate, a black person's chance of unemployment increases by about four percentage points more than a white person's. The same is not true for women, for who we find only weak or no evidence of higher volatility. Through the lens of the model, these findings are consistent with the results of resume studies which tend to show strong evidence for discrimination on the basis of race or ethnicity but no conclusive evidence for hiring discrimination on the basis of gender.

[2] Like age, immigration background, sexual orientation, parenthood, military status and many more. See, for instance, Dahl and Knepper (2023) for a recent discussion in the case of age or Baert (2017) for a more general review.

A central contribution of this paper is to provide some quantitative insights into the effects of hiring discrimination on an individual's labor market outcomes (e.g., unemployment volatility). Specifically, we contribute with a novel mapping from the degree of hiring discrimination to the cyclicality of labor market outcomes . While the degree of discrimination is at least in principle observable, the portion of labor market outcomes that are due to discrimination is not.[3] Leveraging the structure of the standard search-and-matching model of labor market, we provide a way to infer the latter from the former. This mapping is interesting in at least three ways: First, it establishes that there is an extra welfare burden for discriminated groups and provides information to quantify it. For example, a group facing hiring discrimination will obviously have lower average employment. But the higher business cycle volatility means that this group's employment will decrease particularly strongly in recessions, which is when it is particularly painful to not have a job. Second, it allows us to assess counterfactuals. For example, if we can cut hiring discrimination in half, how much higher will the group's employment be in the next recession? Third, as mentioned, resume studies in their basic form can technically only detect "callback discrimination". We show that if there is, in fact, hiring discrimination, we would expect it to show up in differential unemployment volatility, and evidence for such higher business cycle volatility, therefore, gives us an additional data moment consistent with hiring discrimination (although of course, we cannot rule out other potential causes for differential volatility). The paper makes two additional contributions: We extend Blanchard-Diamond's urn-ball model in a tractable way to allow for an arbitrary degree of hiring discrimination. Finally, we also contribute to the relatively thin literature on incorporating racial and gender heterogeneity into the structure of a model focusing on aggregate outcomes.

In summary, this study not only underlines the potential effects of hiring discrimination but also offers a novel framework for understanding its broader economic consequences, particularly over the business cycle. Our results about differential volatility in labor market outcomes adversely affecting discriminated workers should provide an additional frame of discussion for policymakers acting against (undesired) discriminatory practices. The rest of the paper is organized as follows. Section 2 reviews some related literature. Section 3 shows some empirical labor market differences by race and sex using CPS data. Section 4 lays out the basic mechanism of hiring discrimination formalized by an urn-ball matching function and incorporates it into a search model of the labor market, illustrating how it leads to cyclical differences in labor market outcomes. In section 5, we use empirical findings to calibrate the model and quantitatively assess its implications for the labor market impacts of hiring discrimination over the business cycle. Finally, section 6 concludes.

[3]It is observable in the sense that it can be identified in an (idealized) experiment since the unit of observation is an individual. In contrast, one cannot possibly run such idealized experiments on a macroeconomic level.

2 Related Literature

One of the strands of literature this paper is related to is the search/matching theoretical literature that focuses on group differences and heterogeneity. Blanchard and Diamond (1994) use a special case of the urn-ball matching function with lexicographic employer preferences to consider discrimination against long-term unemployed workers. In contrast to their paper, we generalize the matching function to include a continuous margin of discrimination. In their setup, workers become less attractive to employers the longer they remain unemployed; that is, membership in a discriminated group changes over time, endogenizing negative duration dependence of unemployment exit rates for an individual. Our paper studies the cyclical implications of a fixed membership in a discriminated group.

Survey articles by Lang and Lehmann (2012) and Fang and Moro (2011) review work that has focused on the theory of explaining discrimination, in particular with respect to race and gender. These papers, some of which also employ a search-and-matching framework, tend to focus on a possible origin of discriminatory behavior (like taste-based vs information-based) and compare the model implications to differences in average outcomes, like wage or employment gaps. In contrast we are agnostic about the type of discrimination and, taking the rate of discrimination as given, we consider its cyclical effects. Seminal papers in this area are Black (1995), Coate and Loury (1993), and Rosén (1997). Black (1995) shows that if a fraction of employers are discriminatory (they face a utility cost of hiring a minority worker) a wage gap emerges. Coate and Loury (1993) and Rosén (1997) both develop models of statistical discrimination and highlight the potentially self-fulfilling nature of employer beliefs which can operate through incentives for investment in human capital, or through incomplete information about match-specific productivity, respectively.

Another related strand is empirical work on the business cycle differences between groups. Cajner et al. (2017) use CPS data to investigate and decompose racial differences in labor market outcomes, both in regard to levels and volatility. Hoynes et al. (2012) focus on job losses during the 2008/2009 recession and how they were distributed among demographic groups. In contrast to these papers our goal is narrower in that we aim to study specifically the differences in volatilities of unemployment, non-employment and job-finding rates by race and gender and compare those values to our calibrated model. Couch and Fairlie (2010) investigate a "Last Hired, First Fired?" hypothesis for blacks in the US labor market. They do not find that blacks' job-finding rates increase more strongly than whites' during an expansion, a result which we also obtain in our empirical part. As we show below in the model, hiring discrimination does not require differences in the volatility of job-finding rates in order to generate differences in unemployment volatility. The reason is that the effects of job-finding rates on unemployment are non-linear and average levels of job-finding rates differ strongly between blacks and whites.

Our research of course relies on large body of empirical literature on discrimination, of which resume and audit studies constitute a big part. Resume studies in particular, where fictitious applications are submitted to real-world job advertisements, have received renewed interest since Bertrand and Mullainathan (2004). Methodology and relevant findings of these types of experiments are reviewed in Bertrand and Duflo (2016), Neumark (2018), Neumark et al. (2019) or Burn et al. (2022). Baert (2017) presents a collection of correspondence experiments since Bertrand and Mullainathan (2004).

As a whole, this body of research tends to find significant evidence for ethnic and racial discrimination, but considerably less evidence for hiring discrimination on the basis of gender. For example, of the resume studies collected in Baert (2017) that focus on race or ethnicity, only two of 36 fail to find significantly negative effects for minority candidates. Specifically for the situation of blacks in the US labor market, Baert (2017) lists six studies that compare callbacks for applicants with African-American sounding names to such with Anglo-Saxon sounding names. All of those studies find worse response rates for the African-American names with discrimination ratios ranging from 1.16 to 1.50.[4] In section 5 we calibrate our baseline to the median discrimination ratio of those studies (1.38).

In contrast, the situation is not nearly as clear regarding gender discrimination, as is also emphasized by Bertrand and Duflo (2016) and Neumark (2018). There are fewer studies of which a much higher share does not find significant evidence for discrimination against women. Again just counting individual studies listed in Baert (2017) focusing on female versus male applicants' job chances, only two out of eleven find statistically significant levels of discrimination against women, whereas four studies find discrimination against men (and the remaining five studies estimate discrimination ratios not significantly different from 1). There may be some evidence that women are discriminated against when it comes to hiring for occupations that require higher skill levels, are higher paid, or that are traditionally male-dominated (see Riach and Rich, 2002; Neumark et al., 1996), but no systematic picture emerges from the full set of correspondence studies. On the other hand there is at least as much evidence that, vice versa, males are less desired by employers in historically female-dominated jobs or even in sex-integrated occupations (for example in Carlsson, 2011; Booth and Leigh, 2010). Clearly, these findings do not rule out that there are other forms of discrimination against women[5], for example regarding promotions, compensation levels, assignment to tasks and recognition for completed tasks, training, etc. But for the hiring margin we conclude that there is no strong evidence for discrimination on the basis of gender.

For us, resume studies provide a convenient point of comparison in the sense that we can directly compare their estimated callback rate differentials to our parameter of hiring rate differentials. There are, however, two main pieces of information that resume studies cannot identify in their standard design (which most existing studies follow). First, while resume studies can provide clear evidence of discrimination in the callback stage of the hiring process, they do not inform about the effect of group membership on the ultimate hiring decision. The conditional

[4]Specifically, these studies are (discrimination ratios of the respective main specifications in parentheses) Agan and Starr (2017) (1.23), Bertrand and Mullainathan (2004) (1.49), Decker et al. (2015) (1.31), Michael Gaddis (2015) (1.50), Jacquemet and Yannelis (2012a) (1.46), and Nunley et al. (2014) (1.16).

[5]Even in alternative contexts (e.g., Bach et al., 2023; Gharehgozli and Atal, 2020).

hiring rate for an applicant who has passed the callback stage despite being part of a discriminated group could plausibly be greater or smaller than for an applicant of a non-discriminated group, and hence the degree of discrimination could be stronger or weaker than the effects measured by resume studies. However, we think that the effect size measures in these studies is informative at least about the order of magnitude of discrimination for a given group. This issue is also discussed in Neumark (2018) and Riach and Rich (2002), who point out that there are some smaller audit studies finding that most discrimination occurs at the callback rather than the interview stage, and that hence the callback margin may be the most relevant one to study. But it is worth keeping in mind that the relationship between callback and ultimate hiring propensities is not settled empirically. A second issue is that the standard design of resume studies can detect the existence of discrimination, it cannot easily inform about its underlying type: Discrimination may be preference-based or statistical (or both).[6] In the present paper we are correspondingly agnostic about the nature of discrimination.

On the other hand, while there are many studies establishing an average level of discrimination, relatively few of them investigate how the effects of discrimination change over the cycle, at least for race and gender[7]. Baert et al. (2015) find that in the Belgian youth labor market, candidates with foreign sounding names do not receive significantly fewer callbacks during a tight labor market, but do worse than candidates with native sounding names when the labor market is slack. Dahl and Knepper (2023), provide recent evidence of age-related disparities in hiring and firing rates throughout the business cycle. Lastly, Pooling data from earlier studies in Sweden, however, Carlsson et al. (2018) do not find a significant decrease of minority candidates' callbacks in slack labor markets.

Finally, our research contributes to the relatively small yet growing body of literature examining the intersection of macroeconomics and social issues, particularly focusing on differential outcomes across racial groups. Studies such as those by De et al. (2021) and Bartscher et al. (2022) investigate the impacts of monetary policy and macroeconomic shocks on racial groups. Hegarty (2023) explores the role of firm heterogeneity in racial disparities in labor outcomes, while ongoing research by Cairó and Lipton (2023) and Lahcen et al. (2023) delves into the combined effects of discrimination and monetary policy on racially diverse workers. Unlike these studies, our research takes a focused approach on the cyclical effects of hiring discrimination, introducing a distinct element of heterogeneity—hiring discrimination—into a search model using a novel, modified urn-ball matching function. Our model, anchored in empirical resume studies, directly integrates the relative likelihood of hiring discrimination, thus providing a strong empirical foundation. This enhances its relevance to real-world scenarios and offers a clearer understanding of the policy implications of hiring discrimination.

[6] There are some studies that try to disentangle the two in addition to experimental work (see the survey in Bertrand and Duflo, 2016).

[7] There is some evidence following the seminal paper by Kroft et al. (2013) that discrimination by unemployment duration becomes weaker in recessions, consistent with statistical discrimination (in recessions, unemployment duration is a weaker signal of applicant quality).

3 Empirics

This section utilizes the Current Population Survey (CPS) to evaluate the relative volatility of unemployment rates among different groups throughout the business cycle.[8] While the results might appear familiar in the literature, being comparable to those found in Cajner et al. (2017) and Hoynes et al. (2012), conducting this exercise is crucial. It not only underscores differences in labor market outcomes across groups but also provides an empirical framework to refine the proposed model.

The CPS employs a 4-8-4 rolling panel survey pattern for housing units: Residents are included in the dataset for four consecutive months, exit for the subsequent eight months, and then re-enter for another four months. A notable limitation of this dataset, stemming from its focus on housing units rather than individual residents, is the potential for sample selection bias. For example, if a previously unemployed person relocates due to new employment and consequently exits the sample, it could disproportionately affect certain groups. Within each household, all individuals are surveyed.

We concentrate our analysis on women and blacks for several reasons related to our research design. Firstly, numerous resume studies have explored potential biases in the hiring processes affecting these groups. Secondly, gender and race are empirically specific characteristics: they are consistently measured in the CPS (unlike, for example, sexual orientation), and their group memberships are relatively well-defined. For most individuals, these characteristics are binary and stable, unlike other variables such as immigration background or disability status. Finally, gender and race are not endogenous to labor market conditions or employers' hiring practices, unlike factors such as long-term unemployment or parenthood[9].

We utilize monthly data spanning from 1984 to 2018, excluding individuals younger than 25 or older than 55 years, retirees, and members of the armed forces.[10] As a measure of the state of the business cycle we use the aggregate unemployment rate as provided by the the Bureau of Labor Statistics (BLS). The use of the aggregate unemployment rate as a proxy for the business cycle aligns with standard practices, drawing upon the empirical relationship established by Okun's Law,[11] the NBER's definition of a U.S. recession (which encompasses both unemployment and

[8] Data were obtained from IPUMS USA (Flood et al., 2023).

[9] This latter point is helpful because we are interested in the business cycle effects of a constant degree of discrimination. In contrast, for example, long-term unemployment naturally has a higher incidence in recessions which in turn may lead employers to change their behavior towards long-term unemployed applicants (e.g. Jarosch and Pilossoph, 2018, provide evidence that in recessions employers discriminate against long-term unemployed to a lesser degree, as is consistent with statistical discrimination).

[10] This approach aligns with the empirical literature focusing on the 'prime working age' group, often used as a more accurate indicator of the labor market's general condition. It omits younger individuals who may frequently transition between education and employment, and older individuals who are likely retired. As a robustness check, we examined how variations in age criteria might influence our central conclusions and found no significant impact (see Appendix D).

[11] For a recent empirical discussion about the Okun coefficients, see, for instance, Bod'a and Považanová (2021).

output), and the focus on the labor market and human capital aspects of the economy.

For our baseline analysis, we employ pooled OLS, with an unemployment dummy as the outcome variable, which is assigned a value of 1 if the individual is jobless but seeking employment. In our regression, we concentrate on the coefficient of the interaction between the sex/race dummy variable and the unemployment rate. This coefficient will reveal the extent to which the group-specific likelihood of employment (conditional on several demographic characteristics described below) changes relative to a white male's likelihood of unemployment for each one percentage point increase in the aggregate unemployment rate.

While we cannot perfectly replicate the setup of resume studies with observational data, where gender or race is assigned randomly and all other factors are held constant, we can control for a significant number of observable demographic characteristics available in the CPS data. Importantly, these variables are generally observable by employers during the hiring process, ensuring that we are not conditioning on factors that employers cannot observe.[12] Our vector of regressors includes individual characteristics such as age, age-squared, educational attainment[13], and family status[14]. Additionally, we include variables like city size, state, metropolitan area, or occupation to capture fundamental differences across labor markets (e.g., the variability in labor markets across different industries, states, or metropolitan areas in the U.S.). The occupation variable reflects the individual's current employment or, in the case of unemployment, their most recent job. Due to our interest in the interaction terms of gender/race with the state of the business cycle, all controls are included both in levels and as interaction with the aggregate unemployment rate.

The linear probability model in equation (1) summarizes our main regression specification for a person i at time t:

$$y_{it} = \gamma_1 \text{black}_{it} + \gamma_2 (u_t \cdot \text{black}_{it}) + \gamma_3 \text{female}_{it} + \gamma_4 (u_t \cdot \text{female}_{it}) + \boldsymbol{x}'_{it}\boldsymbol{\beta}_1 + u_t \cdot \boldsymbol{x}'_{it}\boldsymbol{\beta}_2 + \varepsilon_{it} \quad (1)$$

where y is the unemployment dummy, $\boldsymbol{x}$ is a vector of controls (including a constant), u the state of the labor market as measured by the aggregate unemployment rate, and $black$ and $female$ are dummies for membership in the respective demographic groups as defined by the CPS.

Because equation (1) assumes that the black-white difference in the expected outcome value, conditional on X, is consistent across genders, we also explore potential sex-related heterogeneity through additional robustness checks, introducing further interactions between race, sex, and

[12]The selection of variables also adheres to standard empirical practices in labor economics research, drawing on frameworks like the human capital Ben-Porath model or the Mincer earnings function. For example, the inclusion of age and age-squared is intended to capture the diminishing returns to experience. As suggested by Forsythe and Wu (2021), using CPS data, using age and potential experience yields similar results.

[13]Categories include less than high school, high school, some college, college, and post-graduate degree.

[14]Categories include married without children, married with children, unmarried without children, and unmarried with children.

the aggregate unemployment rate (see Appendix D). Our findings indicate that the inclusion of these variables minimally impacts the central conclusions drawn from the coefficient of the race-unemployment interaction (Black x Unemployment).

Table 1 presents the estimated coefficients for race and gender, specifically γ_1 through γ_4 from equation (1). Column 1 shows the main specification we previously described, focusing on the interactions between group status and the aggregate unemployment rate. The coefficient for 'Black $\times$ Unemployment,' approximately 0.0039, suggests that (on average) for every percentage point increase in aggregate unemployment, the unemployment rate for blacks rises by about 0.4 percentage points more than that for whites. For instance, during a severe recession where aggregate unemployment escalates from five to ten percent (such as during the Global Financial Crisis), the unemployment gap for blacks relative to whites, using average values for other control variables, would sharply increase from 2 to 4 percent, which is an economically large effect. Conversely, we generally observe neither statistically nor economically significant effects for women. This outcome aligns with resume study findings, which show notable differences in callback rates for blacks but only marginal (if any) differences for women.

The observed pattern is consistent across a range of robustness checks, detailed in the subsequent columns of Table 1. These checks include incorporating state dummies, employing alternative industry controls, and utilizing state-level unemployment rates in place of the aggregate unemployment rate. Furthermore, acknowledging the strong cyclical nature of labor force participation, we further tested the robustness of our results by using the non-employment rate (which includes both unemployed individuals and those not in the labor force) as an alternative outcome measure (see Appendix D). Across these variations, we consistently observe a very similar pattern.

As we will discuss, in the search-and-matching model we propose, the primary mechanism driving differential unemployment rates is attributed to variations in hiring rates. Accordingly, we also examined the empirical behavior of job-finding rates throughout the business cycle. As illustrated in Table 2, the disparity in the probability of finding a job out of unemployment between blacks and whites does not significantly increase during a recession. In fact, the coefficient on the interaction is positive, indicating that this gap narrows as aggregate unemployment rises. However, a significant difference is observed in the baseline levels of job-finding rates, as evidenced by the substantial negative coefficient on the Black dummy. This result will be consistent with our theoretical model: Due to the lower baseline level of job-finding rates for blacks, smaller fluctuations can have a relatively larger impact on their absolute unemployment numbers.

Table 1: Unemployment Status and Business Cycle.

	Baseline	State Dummies	Time-trend	State and Time FE	Industry Dummies	State Unemployment
Black	0.0018*	0.0027**	0.0016	0.0027**	0.0043***	0.0028***
	(2.06)	(2.92)	(1.72)	(2.89)	(4.62)	(3.63)
Black × Unemployment	0.0039***	0.0040***	0.0039***	0.0040***	0.0046***	0.0036***
	(25.97)	(25.6)	(25.52)	(25.55)	(29.27)	(28.04)
Female	0.0055***	0.0056***	0.0055***	0.0056***	0.0058***	0.0042***
	(10.52)	(10.37)	(10.21)	(10.43)	(11.48)	(9.32)
Female × Unemployment	0.00008	0.00009	0.00007	0.00009	0.0001	0.0003**
	(0.90)	(0.93)	(0.81)	(0.95)	(1.19)	(4.19)
R^2	0.041	0.043	0.042	0.043	0.037	0.044
Observations	18,567,096	18,567,096	18,567,096	18,567,096	18,567,096	18,567,096

Notes: 1. This table presents the results for equation (1) using CPS data over 1981m1–2018m12. 2. The dependent variable is a dummy indicating unemployment status, $y = \mathbb{1}(\text{unemployed})$. 3. The first column displays results for the $female$ and $race$ variables in equation (1). 4. The second column includes results after adding state dummies to the control set ($\boldsymbol{x}$). 5. Columns 3, 4, and 5 present results with the addition of a time trend, state and time fixed effects, and 2-digit industry dummies, respectively. 6. The final column shows results using the state-level unemployment rate as the business cycle indicator (u). 7. Survey weights available in the CPS were used in the regression analysis. 8. T-statistics, calculated with robust standard errors, are shown in parentheses. 9. Key: *** significant at 1%; ** significant at 5%; * significant at 10%.

Table 2: Job-finding Rate and Business Cycle

	Job-finding Rate
Black	-0.0637***
	(-10.67)
Black × Unemployment	0.0046***
	(5.32)
Female	-0.0314***
	(-6.08)
Female × Unemployment	0.0006
	(0.74)
Observations	612,034

Notes: 1. This table presents the results for job-finding rates, defined as the rate at which individuals change their unemployment status. 2. The column displays results for the *female* and *race* variables. 3. Results from other control variables ($\boldsymbol{x}$) are omitted for brevity. 4. Survey weights available in the CPS were used in the regression analysis. 5. T-statistics, calculated with robust standard errors, are shown in parentheses. 6. Key: *** significant at 1%; ** significant at 5%; * significant at 10%.

4 Model

In researching the business cycle effects of discriminatory hiring, a challenge is the inability to observe potential outcomes in scenarios other than those that actually occur. Consequently, our research design utilizes a macroeconomic model conducive to counterfactual experimentation. We chose the most standard labor model in macroeconomics for its ease of understanding and its widely accepted predictions. Thus, into this model, we incorporate a source of heterogeneity at the hiring margin. This approach provides a mapping from the degree of hiring discrimination to the cyclicality of labor market outcomes.

Notwithstanding, we acknowledge that, in reality, numerous sources of heterogeneity can influence labor market outcomes. Our focus on the hiring margin does not preclude the existence of other types of discrimination or their interplay with various policies. Factors such as race, age, experience, sex, and educational attainment, for example, are linked to productivity differences or employer preferences. These factors may lead to varying hiring or firing probabilities, resulting in diverse labor market outcomes. Additionally, macroeconomic shocks or policies might affect labor market outcomes differently across population groups. For example, studies like De et al. (2021) and Bartscher et al. (2022) investigate the differential impacts of monetary policy or macroeconomic shocks on African-American workers. Nevertheless, our primary goal is to address the fundamental question of labor outcome volatility when workers face different hiring probabilities. Thus, before exploring more complex scenarios, it is essential to establish a clear understanding of this crucial aspect.

Thus, the model is a search-and-matching model that incorporates an urn-ball matching technology, drawing inspiration from Blanchard and Diamond (1994). Central to our model is a matching mechanism that encapsulates the competition among workers and the varied preferences of employers for different worker types. We have expanded upon Blanchard and Diamond (1994)'s framework to encompass a range of discrimination degrees. Their model posits that when two workers from distinct groups compete for the same vacancy, the worker from the preferred group is invariably chosen. Conversely, our model introduces a scenario where, if two workers from different groups are in the same applicant pool, the worker from the preferred group has an increased, yet not absolute, likelihood of being selected. Consequently, we define the degree of discrimination as the relative hiring probability between two candidates, contingent upon their presence in the same applicant pool. In other words, the relative likelihood of being hired for two otherwise identical workers. Thus, we will assign this relative likelihood to a key parameter in the model. Notably, this parameter mirrors the focus of resume studies but with a distinct emphasis: while resume studies primarily observe relative callback probabilities, our model emphasizes the (relative) hiring probabilities.

4.1 Environment

We briefly describe the economic environment, most of which we keep standard, before explaining the matching function in the following subsection.

Time progresses in discrete intervals. The workforce comprises a unit mass of workers, divided into two demographic groups: N_1 in group 1 and $N_2 = 1 - N_1$ in group 2. The only source of heterogeneity among these groups is their membership; in particular, there are no productivity disparities between workers. A worker from group i (where $i = 1, 2$) earns a wage w_i, and the process of wage determination is detailed in section 4.3. Unemployed workers receive a benefit b. All unemployed workers look for jobs, and their probability of matching with an employer is given by p_i, which is determined by the matching technology described below. Employed workers retain their positions until they are separated from their jobs with an exogenous probability s. In a subsequent section, we will explore the effects of endogenous separations, but for the moment, our emphasis is on the hiring margin while aligning other aspects of the model with the standard framework. Workers consume their entire income within the same period, are risk-neutral, and have a discount factor of β.

Firms are also risk-neutral and share the same discount factor. They maximize expected profits by deciding whether to open a new job vacancy. There is a large number of potential entrants to whom entry into the market is free, such that in equilibrium, the total number of active firms will be determined by a zero-profit condition. Maintaining an open vacancy incurs a cost of c per period. The likelihood of a firm matching with a job seeker is represented by q_i, a value determined by the matching technology as an equilibrium object. Upon successfully matching and hiring a worker, a firm produces a constant output y for the duration of the employment.

This output level is uniform across all firms and workers but fluctuates stochastically over time, thus driving the business cycle dynamics.

It is noteworthy that while the aggregate unemployment rate serves as a proxy for the business cycle in the empirical section, the theoretical model identifies output per match as the key variable for describing economic fluctuations. Nonetheless, there is a clear connection between these two variables. Intuitively, the output per match reflects labor productivity, which directly influences firms' decisions to create or withdraw job vacancies. As a result, higher output per match typically promotes a greater rate of job creation, facilitating employment opportunities for unemployed workers and thus contributing to a reduction in aggregate unemployment.[15]

4.2 Matching function

The matching technology in our model determines the job-finding probabilities, p_i, and the worker-finding probabilities, q_i, as functions of the number of job-seekers and vacancies. We use an urn-ball matching technology as an intuitive way to model the search frictions in the labor market. In this setup, each application submitted by a worker is represented as a ball, and each vacancy as an urn. Every period, during the application stage, each unemployed worker submits one application at random to one of the posted job openings – figuratively, every ball is randomly placed in one of the available urns. If there are many urns and balls, the law of large numbers ensures a fixed distribution of balls across urns. This means that there will be a certain fraction of urns with zero balls, another fraction with exactly one ball, and so forth. Once all applications have been assigned to employers in this way, all employers who have received at least one application hire one of the applicants by drawing one ball out of the respective urn.

We, therefore, assume that an employer will pick between applicants of the same group with equal probability. However, on average, employers exhibit a bias against one of the two groups. For the sake of illustration, let's assume they favor workers from group 1 over group 2. We capture this bias through the parameter π, which represents the relative likelihood of a group 1 worker being chosen over a group 2 worker for a job. For instance, if $\pi = 2$, in an applicant pool containing one worker from each group, the group 1 worker's chance of getting the job is twice that of the group 2 worker, regardless of the size and composition of the rest of the applicant pool. To put it simply, in a scenario where these two are the only candidates, this would mean their respective hiring probabilities are $2/3$ for the group 1 worker and $1/3$ for the group 2 worker.

We consider the parameter π as given and constant over the business cycle. This can be interpreted as a manifestation of taste-based discrimination. Imagine employers making a logit-type choice among applicants, based on a latent, match-specific random variable reflecting the hiring manager's personal preference, which is unrelated to the applicant's productivity. However, this latent preference variable might be correlated with the applicant's group membership. We also

[15] Related literature more closely examines the empirical relationship between vacancies and unemployment. See, for instance, Leythienne (2023).

have a conjecture that the same mechanism works when discrimination is statistical in nature, although to formally show this would require changing the model setup to allow for productivity differences.[16]

In the context of the urn-ball model, applications from different workers are represented by balls of different types, for example, red and white. Formally, let Ω represent the number of urns and Υ the number of balls; where Υ_1 are the type 1 balls (reds), and Υ_2 are the type 2 balls (whites). Each ball is placed in an urn at random with uniform probability across all urns. We define the ratio of balls to urns as the market tightness $\theta = \frac{\Omega}{\Upsilon}$, with $\theta_1 = \frac{\Omega}{\Upsilon_1}$ and $\theta_2 = \frac{\Omega}{\Upsilon_2}$ representing the market tightness for each type of ball, respectively. Since all balls are placed independently, the number of balls in any given urn follows a binomial distribution. Consequently, if both Ω and Υ are large, this distribution can be approximated by a Poisson distribution with the parameter $1/\theta$. In that case, the probability for an individual urn to have k balls placed in it is:

$$\Pr(k;\theta) = \frac{e^{-\frac{1}{\theta}}}{\theta^k k!}$$

Taking into account the different types of colors of balls, because all balls are distributed independently, the probability of having k_1 red balls and k_2 white balls is simply the product:

$$\Pr\left(k_1, k_2; \theta_1, \theta_2\right) = \left(\frac{e^{-\frac{1}{\theta_1}}}{\theta_1^{k_1} k_1!}\right)\left(\frac{e^{-\frac{1}{\theta_2}}}{\theta_2^{k_2} k_2!}\right) = \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1}\theta_2^{k_2} k_1! k_2!}$$

noting that $(1/\theta) = (1/\theta_1) + (1/\theta_2)$ from the definition of the market tightnesses. Moreover, by the law of large numbers, the total number of urns with (k_1, k_2) balls in them is then

$$\Omega \Pr\left(k_1, k_2; \theta_1, \theta_2\right) = \Omega \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1}\theta_2^{k_2} k_1! k_2!}$$

Once the assignment of balls to urns has been made, one ball is drawn at random from every urn. The two types of balls have different probabilities of being drawn from a given urn; without loss of generality, let's say that the red balls are (weakly) more likely to be drawn. Let $\pi \geq 1$ represent the relative likelihood that a red ball is picked compared to a white ball. For any given urn in which there are k_1 red and k_2 white balls, the respective probabilities of drawing a red or white ball are given by

$$\begin{aligned}
\Pr_{1|k_1,k_2}\left(k_1, k_2\right) &= \frac{k_1 \cdot \text{size}_1}{k_1 \cdot \text{size}_1 + k_2 \cdot \text{size}_2} = \frac{\pi k_1}{\pi k_1 + k_2} \\
\Pr_{2|k_1,k_2}\left(k_1, k_2\right) &= \frac{k_2}{\pi k_1 + k_2}
\end{aligned}$$

[16] The basis of this conjecture is as follows: Suppose productivity is worker-specific, and each worker's productivity is drawn from a group-specific distribution. In particular, mean productivity can vary between groups; say, group 1's average productivity is higher than group 2's. During the hiring process, employers receive a noisy, independent, and identically distributed (iid) signal about an applicant's quality.

By the law of large numbers, the total number of red and white balls drawn from all urns combined can be estimated as

$$\#reds = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \Pr(k_1, k_2; \theta_1, \theta_2) \cdot \Pr_{1|k_1,k_2} = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{\pi k_1}{\pi k_1 + k_2}$$

$$\#white = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \Pr(k_1, k_2; \theta_1, \theta_2) \cdot \Pr_{2|k_1,k_2} = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{k_2}{\pi k_1 + k_2}$$

Finally, let p_1 and p_2 denote the probabilities for any red and white ball, respectively, to be drawn from an urn. These probabilities are calculated as the number of total balls drawn relative to all balls of the same color:

$$p_1(\theta_1, \theta_2) = \#reds / \Upsilon_1 \quad = \theta_1 \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{\pi k_1}{\pi k_1 + k_2} \tag{2}$$

$$p_2(\theta_1, \theta_2) = \#whites / \Upsilon_2 = \theta_2 \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{k_2}{\pi k_1 + k_2} \tag{3}$$

Returning from the urn-ball analogy to the economic model, $p_i(\theta_1, \theta_2)$ represent the respective job-finding probabilities which households can calculate knowing the number of vacancies and the number of unemployed workers in both groups (i.e., both market tightnesses θ_1 and θ_2). As equations (2) and (3) indicate, these probabilities are determined by the likelihood of competing with k_1, k_2 other applicants for the same job times the likelihood of being selected from that applicant pool, integrated over all possible combinations of applicant pools (k_1, k_2).

Finally, the probability for a firm to find a worker of type i follows as usual as

$$q_i(\theta_1, \theta_2) = \frac{p_i(\theta_1, \theta_2)}{\theta_i}$$

4.3 Wage setting and value functions

To complete our model, we must specify the wage-setting rule. In principle, we are at liberty to choose any rule that shares the joint match surplus (imposing that all matches lead to hires and match probabilities equal transition probabilities from unemployment to employment and vacancy to filled job, respectively). In the illustrative example of the next section and the central calibration in section 5, we will use a standard Nash-bargaining rule. As alternatives, we also explore the role of constant wages and a no-discrimination policy under which wages are required to be equal across groups.

We now summarize the model by outlining the implied value functions. The state of the

For two candidates with identical underlying productivity, a group-2 candidate requires a higher signal to be hired over a group-1 candidate, leading to differential hiring likelihoods from the same pool. We hypothesize that these relative likelihoods remain relatively stable over the business cycle, as long as the group-specific means among unemployed workers do not fluctuate excessively.

economy is characterized by c θ_1 and θ_2, along with exogenous productivity y. The match-finding probabilities and wages are functions of these state variables. Consequently, a worker of type i receives wage income and either continues in the same job or, at the fixed probability s, becomes unemployed. We later relax the assumption of constant separation rates to explore the role of endogenous separations. The associated value function for a worker is given by

$$W_i(\theta_1, \theta_2, y) = w_i(\theta_1, \theta_2, y) + (1-s)\beta\mathbb{E}_y\left[W_i\left(\theta_1', \theta_2', y'\right)\right] + s\beta\mathbb{E}_y\left[U_i\left(\theta_1', \theta_2', y'\right)\right]$$

Unemployed workers receive unemployment benefits and have a chance to find work, otherwise remaining unemployed in the next period:

$$U_i(\theta_1, \theta_2, y) = b + \beta p_i(\theta_1, \theta_2)\mathbb{E}_y\left[W_i\left(\theta_1', \theta_2', y'\right)\right] + \beta\left[1 - p_i(\theta_1, \theta_2)\right]\mathbb{E}_y\left[U_i\left(\theta_1', \theta_2', y'\right)\right]$$

The current period return for firms in a match is the output produced minus the wages paid. Their continuation value is the expectation of staying in the match versus separating from the worker exogenously. For a firm in a match with a worker of type i, the value function is

$$J_i(\theta_1, \theta_2, y) = y - w_i(\theta_1, \theta_2, y) + (1-s)\beta\mathbb{E}_y\left[J_i\left(\theta_1', \theta_2', y'\right)\right] + s\beta\mathbb{E}_y\left[\max\left\{V\left(\theta_1', \theta_2', y'\right), 0\right\}\right]$$

An unfilled posted vacancy costs a firm an amount c per period and is filled with a worker of type i with probability $q_i(\theta)$. The value of a vacancy is, therefore,

$$V(\theta_1, \theta_2, y) = -c + \beta\sum_i q_i(\theta_1, \theta_2)\mathbb{E}_y\left[J_i\left(\theta_1', \theta_2', y'\right)\right] + \left(1 - \sum_i q_i(\theta_1, \theta_2)\right)\beta\mathbb{E}_y\left[\max\left\{V\left(\theta_1', \theta_2', y'\right), 0\right\}\right]$$

As is standard, the assumption of free entry by firms to post vacancies implies zero profits in expectation; that is, $V(\theta_1, \theta_2, y) = 0$ at all times. Appendix A collects all the equilibrium conditions.

5 Calibration and Results

In this section, we discuss the calibration of the model and present its results. Rather than immediately delving into numerical evaluations post-calibration, we initially focus on analyzing the qualitative properties of the model. This approach is intended to build an intuitive understanding of the central mechanism. Subsequently, we present the quantitative results. Firstly, we assess the model's performance by comparing its outcomes with those derived from the empirical analysis using CPS data. Secondly, we utilize the model for counterfactual exercises. Finally, we conduct a series of robustness checks. As indicated by existing literature, variations in the wage setting or the separation process can influence the predictions of the baseline search-and-matching model. Therefore, we investigate how such alterations impact the central results of our study.

5.1 Baseline calibration

We calibrate the model to match the U.S. economy, focusing on aggregate labor market statistics, leaving as free statistics any associated with differences between groups (e.g., differences between blacks and whites); that is, not taking any group differences into account and leaving them to assess later the model's performance in capturing these disparities. Specifically, our calibration targets the long-run average level of the aggregate unemployment rate and the average job-finding rate. We then use the model to examine cyclical differences in labor market outcomes between groups, given a specific level of hiring discrimination, denoted by π. In other words, we look at how the groups' unemployment and job-finding rates respond following aggregate shocks and how large a difference in these responses a given value of π generates.

Table 3 presents the parameter values used in our calibration. For instance, considering the data's monthly frequency, we select $\beta = 0.9967$, corresponding to an annual discount factor of 0.96. In line with Shimer (2005b), we adopt a separation rate of $s = 3.4\%$, and following Hall and Milgrom (2008), we set the flow value of unemployment at $b = 0.71$. This value lies within the typical range of estimates for this parameter. Additionally, in the CPS data, the ratio of blacks to whites among labor force participants is approximately 1 to 9.

Table 3: Calibration

Parameter	Value	Source
β, discount rate,	0.9967	Monthly frequency, annual interest rate of 4%
s, separation rate,	0.0340	Average separation rate (Shimer, 2005b)
b, value of unemployment,	0.7100	Standard value, e.g. Hall and Milgrom (2008)
ρ_y, persistence of productivity,	0.9830	Quarterly autocorrelation of output 0.95
σ_ε, s.d. of productivity innovations,	0.0019	Standard deviation output 1.65%
N_1, pop. share of group 1,	90%	White/Black share in the labor market
ν, employer's bargaining power,	0.585	Calibrated to mean agg UE and JF rate
c, vacancy creation cost,	0.460	Calibrated to mean agg UE and JF rate
π, degree of hiring discrimination,	1.385	Resume studies

Notes: 1. The table reports the parameter values used in the model. 2. Moments used to calibrate ν are the long-run unemployment rate of 6.2% (CPS), and an average job-finding rate of 45% (as in Shimer, 2005b). 3. S.d.: standard deviation. 4. UE: Unemployment. 5. JF: Job finding.

We calibrate both the vacancy creation cost, denoted as c, and the employer's bargaining power, ν, targeting the long-run mean of the unemployment rate in the CPS data at 6.2% and the job-finding rate reported by Shimer (2005a) at 45%. The implied value for the vacancy creation cost is 0.46, a figure commonly used in the literature and corresponding to a vacancy costing approximately 14 days worth of output. Aggregate vacancy creation costs, for which empirical estimates range between 1% to 2%, sometimes serve as an alternative target for c. In our model, the calibrated value for c implies that aggregate vacancy creation costs amount to 1.9% of aggregate output. The calibration yields an employer's bargaining power, $\nu = 0.564$, a value within the commonly used range in the literature, suggesting that employers receive just over half of the joint match surplus. Table 4 shows that the model moments are reasonably close to the targets.

Finally, for our baseline calibration, we select a degree of hiring discrimination, π, of 1.385. This figure represents the median estimate from the resume studies surveyed by Baert (2017), focusing specifically on African-Americans in the US.[17] This value is also close to the point estimate of 1.49 in the seminal study of Bertrand and Mullainathan (2004). It is important to note that π can be determined independently from the other model parameters, as its impact on the model's aggregate behavior is minimal. Instead, its first-order effect is only on the distribution of labor market outcomes between groups.[18]

5.2 Business cycle effects

We now analyze the qualitative properties of the model to build an intuition about the central mechanism. With this goal in mind, we first present the impulse response functions following a positive shock to output per match. Then, we illustrate the results from the model for different values of individual parameters, such as the discrimination rate.

Impulse response functions. Figure 1 displays the response of the model following a shock to output per match y via impulse response functions. Initially, in the pre-shock periods when the economy is in a steady state, with $\pi > 1$, employers are more likely to hire a given group-1 worker than a given group-2 worker from any applicant pool. This disparity is reflected in higher steady-state job-finding rates for group 1 compared to group 2, leading to a relatively higher steady-state unemployment rate among the latter.

As is standard, an increase in output per match causes firms to post more vacancies, resulting in a jump in labor market tightness θ. Hiring exceeds separations for a few periods, causing the unemployment rate to decrease before recovering gradually in a U-shaped manner.

Of primary interest are the differences in unemployment rates and job-finding rates between the two groups, particularly from the perspective of the discriminated-against group 2. Following the positive shock, due to the weakening of competition with workers of group 1, the job-finding rate for workers of group 2 increases particularly strongly. As a result, the unemployment of group 2 decreases by more, narrowing the gap between the two groups' unemployment rates.

[17] Specifically, taking the midpoint of the estimates of Decker et al. (2015) and Jacquemet and Yannelis (2012b) which are 1.31 and 1.46, respectively.

[18] The reason why there is a small effect on aggregate outcomes in the first place is Nash bargaining: Changes in π slightly alter the bargaining position of the workers of different groups when they encounter a new match, which in turn affects wages and the employer's vacancy creation decision. Because these effects work in different directions for the two groups the net effect on aggregates is very small.

Figure 1: Impulse responses to a positive productivity shock

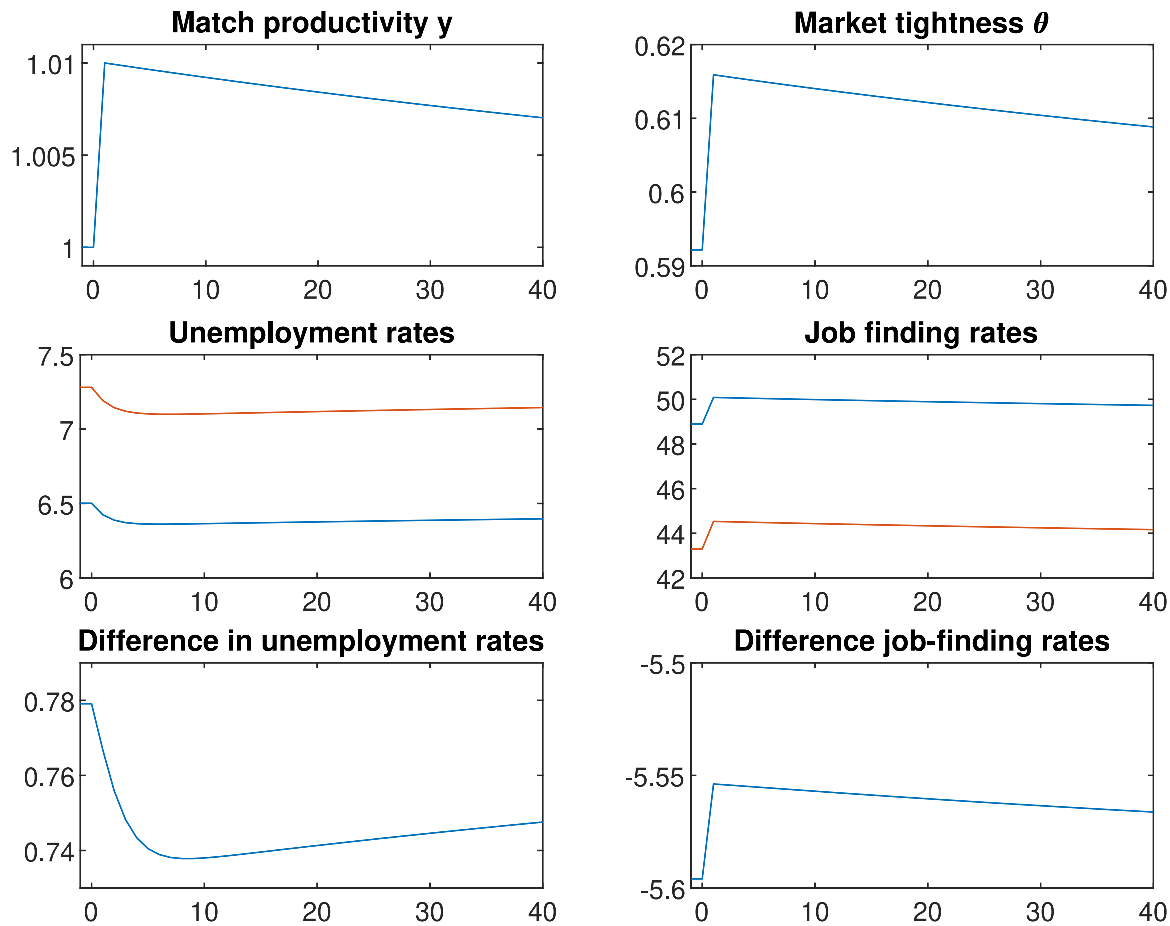

Notes: 1. This figure presents the impulse response functions following a positive shock to output per match. 2. Periods of time are in the x-axis. 3. There are six panels illustrating the temporal evolution of the (i) output; (ii) market tightness; (iii) unemployment rate for each group (blue and red lines represent type-1 and type-2 workers, respectively); (iv) job finding rate for each group; (v) difference in the unemployment rates between the two groups; and (vi) the difference in the job-finding rates between the two groups.

Effect of individual parameters on labor market dynamics. We now investigate how different model parameters affect the resulting differences between the two groups in response to the shock. Specifically, we focus on the parameter of interest π, which determines the relative odds of getting hired for two applicants from different groups in the same applicant pool. Figure 2 displays differences between the groups for different values of π. When $\pi = 1$, workers of both types have an equal chance of being hired from a given pool, resulting in no differences between the groups in the steady-state unemployment or job-finding rates, and thus, identical response to a business cycle shock. For values of π greater than 1, workers of group 1 are hired more readily, to the detriment of group 2 applicants. As described earlier, the discriminated group's unemployment rate responds more strongly to changes in labor market tightness – they are more exposed to the congestion effect of multiple workers applying to the same job posting. The stronger the degree of discrimination, the more pronounced the disparity in the impulse response between groups. However, the effects of increasing π are concave: Even for extremely large values of π, workers from group 2 can find jobs, just as there will be unemployed type-1 workers. In the limit for $\pi \to \infty$, the only chance for a group-2 worker to get hired is to be in a pool without a group-1 applicant, and similarly, a type-1 worker can remain in unemployment if they compete unsuccessfully with one or more applicants of their own group. In this extreme case of discrimination, the model nests Blanchard and Diamond (1994)'s scenario of lexicographic employer preferences.

We also examine variations in the cost of job creation c and the separation rate s (figures 3 and 4 in appendix B). As is standard, both costlier vacancies and shorter expected duration of matches reduce the steady-state labor market tightness θ by making it less attractive for firms to create a new job. Consequently, unemployment is higher and job-finding rates are lower in steady state. However, an increase in these parameters also implies higher volatility in market tightness for a given shock to y, as it raises how much a firm benefits additionally from filling a vacancy. In the model, this greater volatility of θ translates into a larger difference in the groups' response via the congestion mechanism.

5.3 Labor outcome dynamics across worker groups

In this subsection, we present the quantitative results from the model, focusing on central statistics related to aggregate outcomes and outcomes by groups. These results are compared with empirical evidence from CPS data to assess the model's performance.

Quantitative results from the model. Table 4 displays the model's results regarding both aggregate outcomes and outcomes by groups. Most notably, the table shows that there are both strong level effects and strong cyclical effects of hiring discrimination on unemployment. With an unemployment rate of 6.5%, group 2 has about a 0.8 percentage points higher unemployment rate than group 1. Additionally, it takes longer for workers of group 2 to find a job when unemployed, with a substantial gap of 5.6 percentage points in job-finding rates. The cyclical movements in the unemployment rate, a key focus of this paper, are much more pronounced for group 2. The standard deviation of group 2's unemployment rate is about 1.3 percentage points, whereas the standard deviation of group 1's unemployment is 1 – in other words, fluctuations in group 2's unemployment rate exceed group 1's by roughly 30%. It is worth pointing out that these adverse cyclical effects for group 2 do not show up in significantly higher volatility of job-finding rates. This is due to the lower baseline; fluctuations of similar size in job-finding rates constitute larger changes for group 2. The fact that in the data, we do not find large differences in the cyclicality of job-finding rates, but differences in the cyclical behavior of unemployment rates, is consistent with this observation.

Figure 2: Group differences as function of discrimination π (comparative statics)

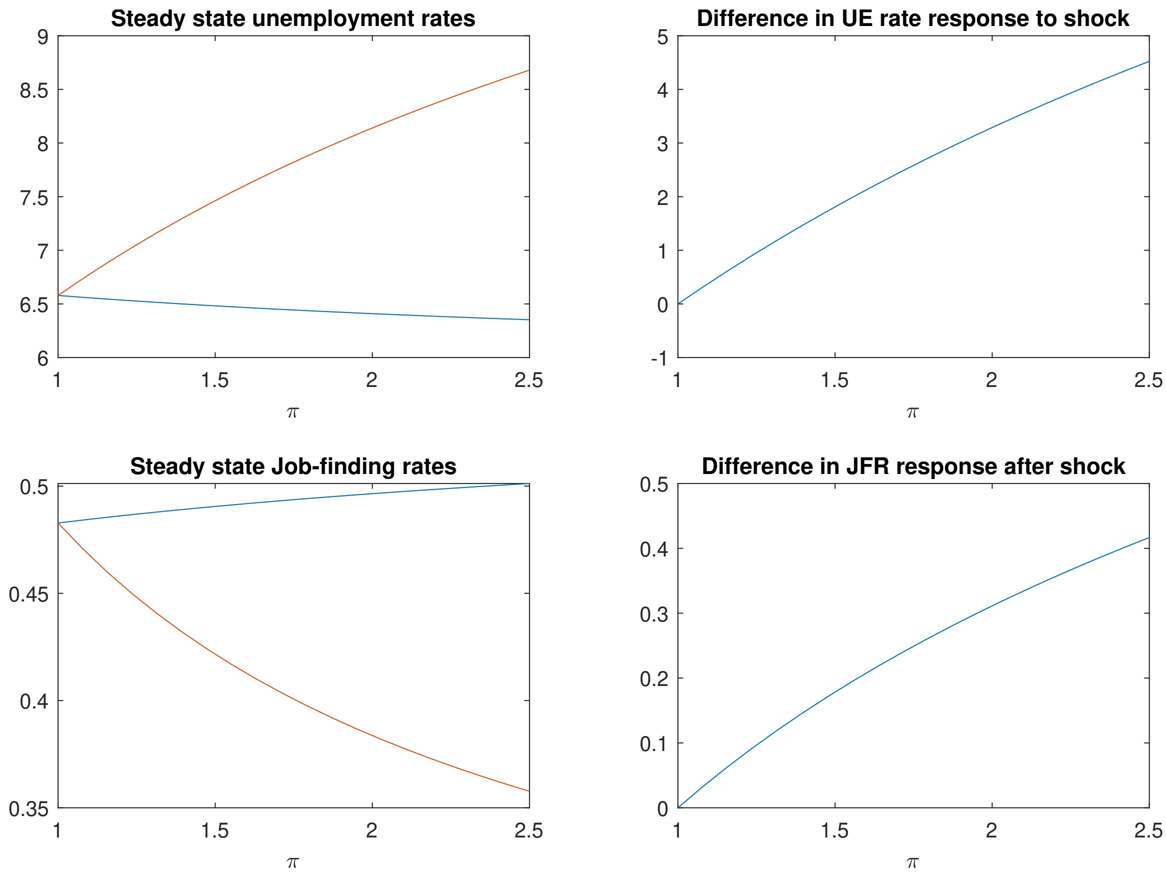

Notes: 1. This figure illustrates the results from the model for different values for the discrimination rate π (x-axis). 2. There are four panels. 3. The left-hand side panels show the steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). 4. The right-hand side panels show the difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity y in unemployment (top right) and job-finding rates (bottom right). 5. Blue and red lines represent the type-1 and type-2 workers, respectively. 6. The asymmetry between groups in the left-hand side panels stems from the calibration in which group 1 comprises a larger share of the population.

To contextualize the numbers in table 4 and compare them with our empirical findings from section 3 using CPS data, consider a numerical example in which aggregate unemployment increases. Based on our point estimate of table 1, in section 3 we had projected that during a severe recession in which average unemployment rises by 5%, the unemployment rate for blacks increases close to 2 percentage points stronger than for whites. The calibrated model implies that in such a recession, the disparity in unemployment rates would increase by 1.39 percentage points, thus accounting for approximately 70% of the empirically measured gap.

Lastly, hiring discrimination does not appear to quantitatively explain the racial wage gap. It is true that, qualitatively, the mechanism generates a difference in wages between groups, and

Table 4: Model results

Outcome	Aggregate	Group 1	Group 2
Mean unemployment	6.54 %	6.47 %	7.23 %
S.d. unemployment	1.05 %	1.02 %	1.32 %
Mean job-finding rate	48.3 %	48.9 %	43.3 %
S.d. job-finding rate	1.69 %	1.69 %	1.75 %
Wages	0.9792	0.9794	0.9771
S.d. wages	0.2145	0.0136	0.0137

Notes: The table reports the steady state values and standard deviations (s.d.) of outcomes in the model.

that the variance of wages is larger for the discriminated group: Because job-finding rates for the minority workers are lower, they have a worse outside option and are hence able to extract less of the output produced during the match. However, as can be seen in the last rows of 4, these wage differences are small. For example, for doing a job that generates \$100 worth of output, a group-1 worker receives compensation of \$97.94, whereas a worker of group 2 gets paid ¢23 less. This gap is, of course, much smaller than the empirically observed racial pay gap as measured by the unexplained component in standard wage regressions. For example, Daly et al. (2017) estimate this gap at around 9% for men and 5% for women.

5.4 Counterfactuals

The model, thanks to the mapping from hiring discrimination to the difference in labor market outcomes, allows us to explore how hiring discrimination impacts labor market outcomes and to conduct counterfactual policy analysis. Two relevant experiments are: (i) How large is the degree of hiring discrimination that we would assume to explain the entire difference in labor market outcomes? and (ii) what is the effect of reducing discrimination? i.e., quantifying how much the gap in labor market outcomes narrows when hiring discrimination is lessened.

The model establishes a map between hiring discrimination and variations in labor market outcomes, which allows us to explore how hiring discrimination impacts labor market outcomes and to conduct counterfactual policy analysis. Two relevant experiments are: (i) What magnitude of hiring discrimination would be required to fully explain the observed differences in labor market outcomes? What would be the impact on these outcomes if the level of discrimination were reduced? In other words, we quantify the reduction in the disparities in labor market outcomes that could be achieved by lessening hiring discrimination."

To address the first question, we increase the value of π until the relative difference in unemployment volatilities aligns with our empirically measured value. This is achieved at a value of $\pi = 1.57$, closely mirroring Bertrand and Mullainathan (2004)'s estimate of 1.49, and is still within the range of other resume studies surveyed by Baert (2017). We can think of this exercise as an attempt to identify the degree of hiring discrimination based on the relative volatility of unemployment rates, conditional on the model being correct. Identifying π in this way implies that, in the model, hiring discrimination now accounts for a steady-state difference in unemployment rates of 1.09 percentage points (compared to 0.85 points in the baseline) and a steady-state difference in job-finding rates of 7.6 percentage points (compared to 5.8 points previously).

Conversely, we can ask how much the additional unemployment of blacks would be reduced if we could, say, cut the amount of hiring discrimination in half. We, therefore, set $\pi = 1.1925$. In the model, this means a reduction of black steady-state unemployment of about 0.33 percentage points, from 7.28% to 6.95%, and an increase in steady-state job-finding rates of about 2.2 percentage points, from 43.3% to 45.5%. The volatility of unemployment is also reduced, now exceeding that of whites by 15.2% compared to 29.7% in the baseline calibration[19]. This reduced volatility implies that in the case of our exemplary big recession with 5% higher aggregate unemployment, the black unemployment rate would increase by 0.73 percentage points more that of whites, as opposed to 1.39 percentage points in the baseline scenario. Thus, a reduction in hiring discrimination could potentially decrease black unemployment by 0.66 percentage points during a severe recession.

5.5 Alternative wage-setting rules

To explore how different wage-setting mechanisms affect our results, we consider two plausible alternatives in the model. The first is a no-discrimination rule, where employers are prohibited from basing wages on group status, and the second is a scenario with constant wages, where wages do not fluctuate with productivity. Thus, while our baseline Nash bargaining rule allows wages to vary based on group status and business cycle conditions, the first alternative eliminates wage variability between groups, and the second removes wage variability over the cycle. This approach enables us to assess the impact of these aspects of the Nash rule separately.

No wage discrimination. Under the first alternative, we assume that wages must be equal across groups at all times. In this scenario, the wage can only depend on productivity, not on group membership, at any point in the business cycle. This setup is inconsistent with standard Nash bargaining, as the two worker groups have different respective outside options (notably, group 1's value of unemployment is higher due to a shorter expected unemployment duration). We, therefore, model the wage as sharing the joint surplus between an employer and the average worker:

$$(1-\nu)\, J = \nu \left[N_1 \left(W_1 - U_1\right) + \left(1 - N_1\right)\left(W_2 - U_2\right)\right]$$

[19] As shown, the difference in labor market outcomes is non-linear in π, although this concavity is not very pronounced in this area of the parameter value.

In this model, there is no distinction for the employer in terms of the value of a match with either group, allowing us to omit the index i on the value of a match for the employer J.

This wage rule results in minimal employment effects as it barely impacts the aggregate wage level (i.e., the cost for employers to post a vacancy). Essentially, the equal-wage requirement leads to redistributive outcomes: wages of group 1 decrease slightly to the benefit of group-2 wages, and there are no meaningful employment effects for either group. Thus, this rule can neutralize the (small) negative wage effects of hiring discrimination, although the direct effect of hiring discrimination on differential employment dynamics continues to persist.

Constant wages. In this scenario, wages remain constant over time and across workers within each group but differ between groups. That is, at any time, workers of group 1 will be paid wage $\overline{w}_1$, and workers of group 2 will be paid $\overline{w}_2$. It seems natural to pick the respective steady-state values from the Nash baseline for $\overline{w}_1$ and $\overline{w}_2$.[20] As highlighted by (e.g. Hall, 2005) rigid wages can amplify the volatility of other labor market outcomes. Therefore, we recalibrate the productivity process to keep the volatility of output at the baseline level, and we also adjust the calibrated parameters c and ν (vacancy creation costs and bargaining power) to target the same aggregate moments as before (mean unemployment rate and job-finding rate).

Table 5 shows that, as expected, removing wage adjustment over the business cycle increases the volatility of unemployment fluctuations. However, this increase occurs fairly evenly across both groups, meaning their relative standard deviations of unemployment remain roughly the same as under Nash wages.

Table 5: Alternative wage-setting rules

Outcome	Baseline (Nash)		No-discrimination		Constant wages	
	Group 1	Group 2	Group 1	Group 2	Group 1	Group 2
Mean wages	0.9794	0.9771	0.9792		0.9794	0.9771
S.d. wages	0.0136	0.0137	0.0136		0.0000	
S.d. unemployment	1.02	1.32	1.02	1.32	1.13	1.46
Add'l group-2 UE in rec.	1.39		1.39		1.35	

Notes: 1. The table reports the steady-state values and standard deviations (s.d.) of outcomes in the model using three alternative rules determining the wage. 2. Row 'Add'l group-2 UE in rec.': How many percentage points more does group 2's unemployment rate increase compared to group 1 when a recession increases aggregate unemployment by 5 percentage points?

[20] We assume that both sides can commit to staying in the match long enough so that we do not have to verify that the wage stays in the interior of the set for which both parties extract positive shares of the surplus.

5.6 Endogenous separations[21]

As mentioned, our primary goal is to analyze the cyclical impact of documented levels of hiring discrimination using a standard labor market model in macroeconomics. However, the literature suggests that some assumptions in the standard search-and-matching model, such as the process of separation rates, may alter its predictions. In particular, fluctuations in separation rates can also drive unemployment volatility. Fujita and Ramey (2009) found that around half of the fluctuations in unemployment can be attributed to variations in the separation rate, and Couch and Fairlie (2010) documented that blacks are more likely to be fired during economic downturns. Consequently, we explore the implications of modifying the baseline assumption of a constant separation rate. Following the modeling approach proposed by Fujita and Ramey (2012), we allow for endogenous separations in our model. We emphasize that, despite this modification, the primary focus remains centered on the hiring margin.

In their search-and-matching model, Fujita and Ramey characterize workers as heterogeneous in terms of job productivity, with only a fraction s of the separation rate being exogenous. A worker-firm match can therefore produce an output level y, defined as the product $z \cdot x$. Here, z represents an aggregated productivity factor, and x denotes match-specific productivity factors. These factors are indexed by $x \in X = \{x_1, x_2, ..., x_M\}$ for $x_1 < x_2 < ... < x_M = x_h$ and $h \geq 2$. New matches start at $x = x_h$, but the value of x may change in subsequent periods. The probability of a switch at the end of a time period is λ, leading to a random draw of x for the next period from the cumulative distribution function $G(x)$. With probability $(1-\lambda)$, x maintains its time period t value into the next time period. The value of a match with a worker of type i at the production stage is then described by the Bellman equation:

$$J_i(z,x) = \max\left\{V(z,x)\,,\, J_i^c(z,x)\right\}$$

where $J_i^c(z,x)$ represents the value for a firm after continuation of the match is chosen,

$$J_i^c(z,x) = zx - w_i(z,x) + \beta\,\mathbb{E}_z\left\{(1-s)\left(\lambda\int_0^{x_h} J_i(z',x)dG(x) + (1-\lambda)J_i(z',x)\right)\right\} + \beta\,\mathbb{E}_z\left\{s\,V(z',x')\right\}$$

and a worker of type i has now the value function:

$$W_i(z,x) = \max\left\{U_i(z,x)\,,\, W_i^c(z,x)\right\}$$

where $W_i^c(z,x)$ represents the value for a worker after continuation of the match is chosen,

$$W_i^c(z,x) = w_i(z,x) + \beta\,\mathbb{E}_z\left\{(1-s)\left(\lambda\int_0^{x_h} W_i(z',x)dG(x) + (1-\lambda)W_i(z',x)\right)\right\} + \beta\,\mathbb{E}_z\left\{s\,U_i(z',x')\right\}$$

Other equations in the model (e.g., the value function for unemployment or a vacancy)

[21]The extension in this subsection was studied in more detail (e.g., the specification with endogenous separation and with/without on-the-job search) in a chapter of the Ph.D. thesis of one of the co-authors.

essentially remain as described in section 4.

Furthermore, we adopt the calibration used by Fujita and Ramey, especially for the recently incorporated parameters. For example, as described in their paper, the idiosyncratic shocks are drawn from a lognormal distribution with $\sigma_x = 0.18$ and the arrival rate of the match-specific productivity shock (probability of changing idiosyncratic productivity) λ is 0.085. These parameters are set based on labor statistics observed in CPS data (i.e., the inflow rate to unemployment of about 2.2 percentage points).

Table 6 shows the results after incorporating the endogenous separation rate into the model. As expected from the literature, there are changes in the outcomes, particularly a reduction in unemployment volatility. Still, we find that the cyclical movements in unemployment rates are stronger for group 2, with fluctuations in group 2's unemployment rate exceeding group 1's by about 14 percentage points.

Table 6: Model results: Constant and endogenous separation rates

Outcome	Baseline Model		Endogenous Separation	
	Group 1	Group 2	Group 1	Group 2
Mean unemployment	6.47 %	7.23 %	6.16 %	6.48 %
S.d. unemployment	1.02 %	1.32 %	0.95 %	1.07 %

Note: The table reports the steady state values and standard deviations (s.d.) of outcomes in two models: the baseline model with a constant separation rate and a specification with endogenous separation.

6 Conclusion

Discrimination in labor markets raises significant moral and efficiency concerns. Empirical evidence from audit studies suggests the existence of racial discrimination during the early stages of the hiring process (i.e., evidence that black workers have lower callback rates for job interviews than white workers). While such practices lead to disparities in labor market outcomes at the levels (e.g., gaps in unemployment rates between Blacks and Whites), their potential effects on labor market volatility, especially during economic downturns, are less understood. This paper researches whether the volatility of labor market outcomes for disadvantaged groups is intensified due to discriminatory practices.

In particular, we employ empirical evidence from correspondence studies and CPS data to examine the potential effects of hiring discriminatory practices on labor market outcomes over the business cycle in the U.S. For this, we extend the urn-ball matching function within a standard search-and-matching model, offering a novel quantitative mapping from the degree of hiring discrimination to differences in labor market outcomes.

After calibrating the model to the U.S. economy, our findings reveal significant disparities: groups subject to higher discrimination, indicated by lower hiring probabilities, also experience greater unemployment volatility. This result underscores the profound, cyclical ramifications of discriminatory practices beyond just the level of unemployment. We also found that the model replicates a substantial portion (around 70%) of the excess unemployment volatility observed among black workers compared to whites.

These insights are crucial, linking micro-level discriminatory practices to macro-level labor market dynamics. They highlight the importance of comprehensive anti-discrimination policies and their enforcement. While initiatives like the U.S. Equal Employment Opportunity Commission (EEOC) and laws such as the Civil Rights Act and the Age Discrimination in Employment Act are vital, the persistent disparities our model highlights point to the need for more effective enforcement and potentially innovative policy approaches to mitigate the cyclicality of labor market outcomes due to discrimination.

Finally, we contribute to the relatively small body of literature exploring the intersection of macroeconomic models and social issues, particularly in the context of discrimination. By elucidating how discriminatory hiring practices affect labor market volatility, our work provides a framework for further exploration of labor market disparities. Future research can build upon this foundation by incorporating other forms of labor market heterogeneity. For instance, examining the role of age, as empirically discussed in Neumark et al. (2019), Burn et al. (2022), and Dahl and Knepper (2023), would be interesting.[22] Such efforts could enhance our understanding of the intricate dynamics that shape employment and economic cycles.

[22]We thank an anonymous referee for this suggestion.

References

Agan, A. and S. Starr (2017). Ban the Box and Racial Discrimination. Quarterly Journal of Economics 133(1), 734–764.

Bach, T. N., T. Le, T. X. Nguyen, and K. Hoang (2023). Gender discrimination, social networks and access to informal finance of vietnamese small and medium enterprises. Economic Analysis and Policy 78, 358–372.

Baert, S. (2017). Hiring Discrimination: An Overview of (Almost) All Correspondence Experiments Since 2005.

Baert, S., B. Cockx, N. Gheyle, and C. Vandamme (2015). Is there less discrimination in occupations where recruitment is difficult? Industrial and Labor Relations Review 68(3), 467–500.

Bartscher, A. K., M. Schularick, M. Kuhn, and P. Wachtel (2022). Monetary policy and racial inequality. Brookings Papers on Economic Activity 2022(1), 1–63.

Bertrand, M. and E. Duflo (2016). Field Experiments on Discrimination. NBER Working Paper Series 22014, 110.

Bertrand, M. and S. Mullainathan (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. American Economic Review 94(4), 991–1013.

Black, D. A. (1995). Discrimination in an Equilibrium Search Model. Journal of Labor Economics 13(2), 309–334.

Blanchard, O. and P. Diamond (1994). Ranking, unemployment duration, and wages. The Review of Economic Studies 61(3), 417–434.

Bod'a, M. and M. Považanová (2021). Output-unemployment asymmetry in okun coefficients for oecd countries. Economic Analysis and Policy 69, 307–323.

Booth, A. and A. Leigh (2010). Do employers discriminate by gender? A field experiment in female-dominated occupations. Economics Letters 107(2), 236–238.

Burn, I., P. Button, L. M. Corella, and D. Neumark (2022). Does ageist language in job ads predict age discrimination in hiring? Journal of Labor Economics 40(3), 613–667.

Cairó, I. and A. Lipton (2023). Labor market discrimination and the racial unemployment gap: Can monetary policy make a difference?

Cajner, T., T. Radler, D. Ratner, and I. Vidangos (2017). Racial Gaps in Labor Market Outcomes in the Last Four Decades and over the Business Cycle. Finance and Economics Discussion Series 2017(071).

Carlsson, M. (2011). Does hiring discrimination cause gender segregation in the Swedish labor market? Feminist Economics 17(3), 71–102.

Carlsson, M., L. Fumarco, and D. O. Rooth (2018). Ethnic discrimination in hiring, labour market tightness and the business cycle - evidence from field experiments. Applied Economics 50(24), 2652–2663.

Coate, S. and G. C. Loury (1993). Will Affirmative-Action Policies Eliminate Negative Stereotypes? The American Economic Review 83(5), 1220–1240.

Couch, K. A. and R. Fairlie (2010). Last hired, first fired? Black-white unemployment and the business cycle. Demography 47(1), 227–47.

Dahl, G. B. and M. Knepper (2023). Age discrimination across the business cycle. American Economic Journal: Economic Policy 15(4), 75–112.

Daly, M., B. Hobijn, and J. H. Pedtke (2017). Disappointing facts about the black-white wage gap. FRBSF Economic Letter 26.

De, K., R. A. Compton, D. C. Giedeman, and G. A. Hoover (2021). Macroeconomic shocks and racial labor market differences. Southern Economic Journal 88(2), 680–704.

Decker, S. H., N. Ortiz, C. Spohn, and E. Hedberg (2015). Criminal stigma, race, and ethnicity: The consequences of imprisonment for employment. Journal of Criminal Justice 43(2), 108–121.

Fang, H. and A. Moro (2011). Theories of statistical discrimination and affirmative action: A survey. Handbook of Social Economics 1(1 B), 133–200.

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, D. Warren, A. Chen, G. Cooper, S. Richards, M. Schouweiler, and M. Westberry (2023). Integrated public use microdata series (ipums), current population survey: Version 11.0 [dataset].

Forsythe, E. and J.-C. Wu (2021). Explaining demographic heterogeneity in cyclical unemployment. Labour Economics 69, 101955.

Fujita, S. and G. Ramey (2009). The cyclicality of separation and job finding rates. International Economic Review 50(2), 415–430.

Fujita, S. and G. Ramey (2012). Exogenous versus endogenous separation. American Economic Journal: Macroeconomics 4(4), 68–93.

Gharehgozli, O. and V. Atal (2020). Revisiting the gender wage gap in the united states. Economic Analysis and Policy 66, 207–216.

Hall, R. E. (2005). Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market. Review of Economics and Statistics 87(3), 408–410.

Hall, R. E. and P. R. Milgrom (2008). The limited influence of unemployment on the wage bargain. American Economic Review 98(4), 1653–1674.

Hegarty, C. (2023). Firm heterogeneity and racial labor market disparities.

Hoynes, H., D. L. Miller, and J. Schaller (2012). Who Suffers During Recessions? Journal of Economic Perspectives 26(3), 27–48.

Jacquemet, N. and C. Yannelis (2012a). Indiscriminate Discrimination: A Correspondence Test for Ethnic Homophily in the Chicago Labor Market. Labour Economics 19(6), 824–832.

Jacquemet, N. and C. Yannelis (2012b). Indiscriminate discrimination: A correspondence test for ethnic homophily in the Chicago labor market. Labour Economics 19(6), 824–832.

Jarosch, G. and L. Pilossoph (2018). Statistical Discrimination and Duration Dependence in the Job Finding Rate. NBER Working Paper (24200), NBER.

Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment. The Quarterly Journal of Economics 128(3), 1123–1167.

Lahcen, M. A., G. Baughman, and H. van Buggenum (2023). Racial unemployment gaps and the disparate impact of the inflation tax.

Lang, K. and J.-Y. K. Lehmann (2012). Racial Discrimination in the Labor Market: Theory and Empirics. Journal of Economic Literature 50(4), 959–1006.

Leythienne, D. (2023). Is euro area at full employment? – a diagnosis from the beveridge curve. Economic Analysis and Policy.

Michael Gaddis, S. (2015). Discrimination in the credential society: An audit study of race and college selectivity in the labor market. Social Forces 93(4), 1451–1459.

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

Neumark, D., R. J. Bank, and K. D. Van Nort (1996). Sex Discrimination in Restaurant Hiring: An Audit Study. Quarterly Journal of Economics 111(3), 915–941.

Neumark, D., I. Burn, and P. Button (2019). Is it harder for older workers to find jobs? new and improved evidence from a field experiment. Journal of Political Economy 127(2), 922–970.

Nunley, J. M., A. Pugh, N. Romero, and R. A. Seals (2014). An Examination of Racial Discrimination in the Labor Market for Recent College Graduates: Estimates from the Field. Berkeley Electronic Journal of Economic Analysis and Policy 5(3), 1093–1125.

Riach, A. P. A. and J. Rich (2002). Field Experiments of Discrimination in the Market Place. The Economic Journal 112(November), 480–518.

Rosén, Å. (1997). An equilibrium search-matching model of discrimination. European Economic Review 41(8), 1589–1613.

Shimer, R. (2005a). The cyclicality of hires, separations, and job-to-job transitions. Review-Federal Reserve Bank of Saint Louis 87(4), 493 – 507.

Shimer, R. (2005b). The Cyclical Behavior of Equilibrium Unemployment and Vacancies. American Economic Review 95(1), 25–49.

Appendix A Equilibrium conditions

Utility value for an employed and unemployed worker of type i, respectively:

$$W_i(\theta_1,\theta_2,y) = w_i(\theta_1,\theta_2,y) + (1-s)\beta E\left[W_i\left(\theta_1',\theta_2',y'\right)\right] + s\beta E\left[U_i\left(\theta_1',\theta_2',y'\right)\right]$$
$$U_i(\theta_1,\theta_2,y) = b + \beta p_i(\theta_1,\theta_2) E\left[W_i\left(\theta_1',\theta_2',y'\right)\right] + \beta\left[1-p_i(\theta_1,\theta_2)\right] E\left[U_i\left(\theta_1',\theta_2',y'\right)\right]$$

Value to firm of an existing match with worker i of a vacancy, respectively:

$$J_i(\theta_1,\theta_2,y) = y - w_i(\theta_1,\theta_2,y) + (1-s)\beta E\left[J\left(\theta_1',\theta_2',y'\right)\right] + s\beta E\left[V\left(\theta_1',\theta_2',y'\right)\right]$$
$$V(\theta_1,\theta_2,y) = -c + \beta\sum_i q_i(\theta_1,\theta_2) E\left[J_i\left(\theta_1',\theta_2',y'\right)\right] + \left[1-q(\theta_1,\theta_2)\right]\beta E\left[V\left(\theta_1',\theta_2',y'\right)\right]$$

Job-finding probability for worker of type i and probability of firm to find worker of type i:

$$p_i(\theta_1,\theta_2) = \theta_i \sum_{k_1=0}^{\infty}\sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1}\theta_2^{k_2}k_1!k_2!}\frac{\pi k_i}{\pi k_1 + k_2}$$
$$q_i(\theta_1,\theta_2) = \frac{p_i(\theta_1,\theta_2)}{\theta_i}$$

Nash bargaining:

$$J_i(\theta_1,\theta_2,y) - V(\theta_1,\theta_2,y) = \frac{\nu}{1-\nu}\left[W_i(\theta_1,\theta_2,y) - U_i(\theta_1,\theta_2,y)\right]$$

Free-entry (determining number of vacancies):

$$V(\theta_1,\theta_2,y) = 0$$

Definition of market tightnesses as vacancies per group-specific job seeker:

$$\theta_i = \text{vacancies}/u_i$$

Evolution of unemployment:

$$u_i = s\left(N_i - u_{i,t-1}\right) + \left(1-p_i(\theta_1,\theta_2)\right)u_{i,t-1}$$

Exogenous process for match productivity:

$$\log y = \rho \log y_{t-1} + \varepsilon$$

Appendix B Effects of parameters c and s

Figure 3: Group differences as function of vacancy creation costs c

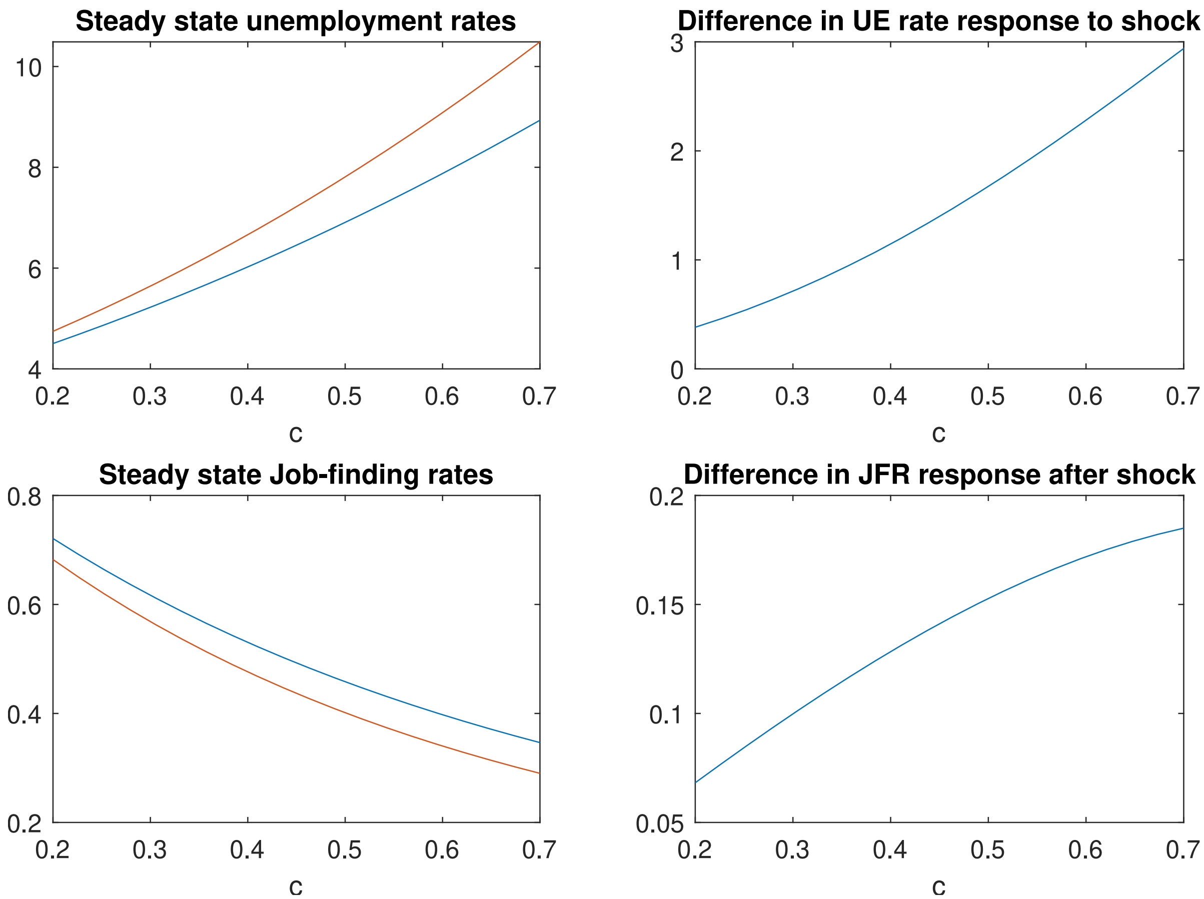

Notes: 1. Model results for different values for the vacancy creation cost c (x-axis). 2. Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). 3. Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity y in unemployment (top right) and job-finding rates (bottom right).

Figure 4: Group differences as function of separation rate s

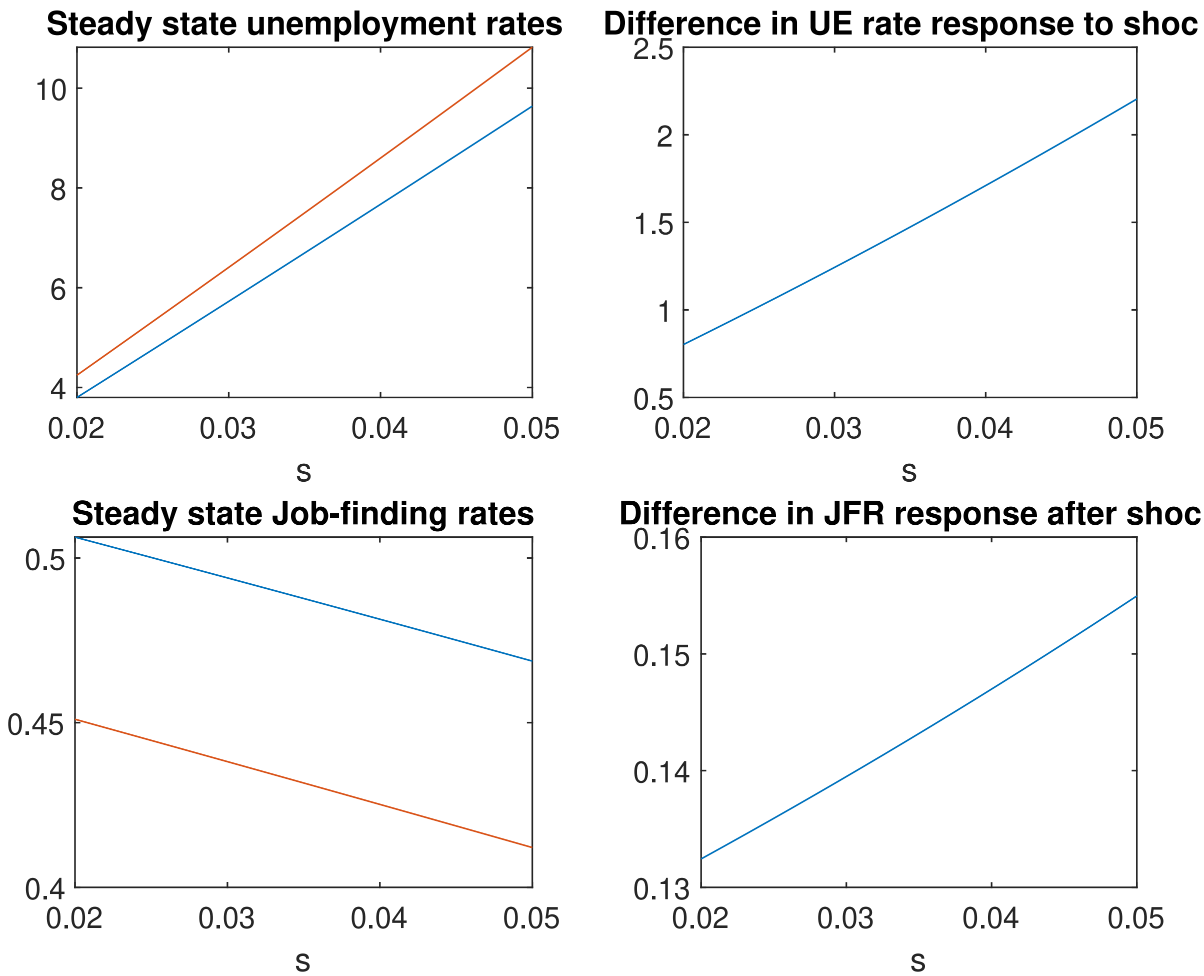

Notes: 1. Model results for different values for the separation rate s (x-axis). 2. Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). 3. Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity y in unemployment (top right) and job-finding rates (bottom right).

Appendix C Descriptive Statistics

Table 7: Descriptive Statistics for Unemployment Rate (%) by Groups.

	Male	Female	Overall
	Mean/(S.d.)	Mean/(S.d.)	Mean/(S.d.)
Panel A. Race			
White	4.11	3.86	4.00
	(19.86)	(19.27)	(19.59)
Black / mixed black	9.45	8.83	9.11
	(29.25)	(28.37)	(28.77)
Hispanic	6.34	7.34	6.77
	(24.37)	(26.07)	(25.12)
Other	5.77	5.27	5.53
	(23.32)	(22.35)	(22.86)
Panel B. Education			
Some high school	9.49	11.08	10.11
	(29.31)	(31.39)	(30.14)
High school or GED	5.95	5.60	5.79
	(23.66)	(23.00)	(23.36)
Some college or associate degree	4.53	4.59	4.56
	(20.79)	(20.93)	(20.86)
Bachelor's degree	2.77	2.92	2.84
	(16.40)	(16.85)	(16.62)
Higher degree	1.99	2.28	2.13
	(13.98)	(14.93)	(14.45)
Panel C. Age			
25 - 30	6.38	6.29	6.34
	(24.44)	(24.27)	(24.36)
31 - 35	5.08	5.36	5.21
	(21.96)	(22.52)	(22.22)
36 - 40	4.57	4.72	4.64
	(20.87)	(21.20)	(21.03)
41 - 45	4.26	4.17	4.22
	(20.19)	(20.00)	(20.10)
46 - 50	4.21	3.90	4.06
	(20.07)	(19.37)	(19.74)
51 - 55	4.18	3.75	3.98
	(20.01)	(19.00)	(19.54)
Overall	4.86	4.79	4.82
	(21.50)	(21.35)	(21.43)
Observations	10,471,501	9,315,910	19,787,411

Notes: 1. This table reports the mean value and the standard deviation (in parentheses) for the unemployment rate (in percentage, %). 2. Author's calculations using CPS data over 1984m1-2018m3. 3. Survey weights available in the CPS data were used for the computation of the statistics.

Table 8: Distribution of individuals by race and sex (%) in the dataset.

	Group							
	Male				Female			
	White	Black mixed black	Hispanic	Other	White	Black mixed black	Hispanic	Other
Panel A. Educational attainment								
Some high school	7,3	2,4	3,0	0,7	4,7	2,5	2,4	0,5
High school/GED	12,7	3,3	1,7	1,1	9,3	3,3	1,3	0,8
Some college	8,3	1,9	0,8	0,8	8,3	2,6	0,9	0,8
Bachelor's degree	4,5	0,6	0,2	0,6	4,4	0,7	0,3	0,6
Higher degree	1,6	0,2	0,1	0,3	1,7	0,3	0,1	0,3
Panel B. Marital status and children								
Married, no children	5,0	0,7	0,6	0,4	4,8	0,6	0,4	0,4
Married, children	12,2	2,1	2,4	1,2	11,3	1,8	1,9	1,1
Non-married, no children	14,9	4,4	1,7	1,2	6,9	2,2	0,7	0,6
Non-married, children	2,6	0,9	0,5	0,3	6,1	4,5	1,3	0,6
Panel C. Age								
25 – 35	13,6	3,5	2,3	1,2	11,0	4,4	1,9	1,1
35 – 45	11,1	2,5	1,6	0,9	9,6	2,8	1,5	0,9
45 – 55	10,2	2,0	1,1	0,9	8,5	1,8	0,9	0,7
Panel D. Industry								
Agriculture (AFF)	1,5	0,3	0,6	0,2	0,5	0,1	0,3	0,0
Mining/construction	9,7	1,4	1,4	0,6	0,9	0,1	0,1	0,1
Manufacturing	6,6	1,5	0,9	0,5	4,2	1,3	0,8	0,4
Transportation/comm.	2,7	0,8	0,3	0,2	1,2	0,4	0,1	0,1
Wholesale trade	1,2	0,2	0,2	0,1	0,8	0,1	0,2	0,1
Retail trade	4,7	1,2	0,6	0,5	6,7	1,8	0,8	0,5
Finance/insurance (FIRE)	1,2	0,3	0,1	0,1	1,9	0,5	0,2	0,2
Services/others	7,9	2,4	0,9	0,9	13,1	4,7	1,7	1,3
Panel E. Overall								
Proportion	36.9	8.2	5.3	3.3	29.5	9.4	4.5	2.9

Note: 1. This table details the distribution of individuals in the CPS data. 2. The columns show information by race and sex. 3. The rows present the information organized according to educational attainment, marital status (with a distinction between those with and without children), age, and industry affiliation. 4. Author's calculations using CPS over 1984m1–2018m12.

Appendix D Additional Empirical Checks

Table 9: Alternative definition of the dependent variable, $y = \mathbb{1}(\text{non-working})$.

	Baseline	State Dummies	Time-trend	Industry Dummies	State Unempl.
Black	0.00153	0.00272**	0.00135	0.00514***	0.00270***
	(1.64)	(2.85)	(1.45)	(5.37)	(3.52)
Black X Unemployment	0.00373***	0.00384***	0.00379***	0.00441***	0.00342***
	(23.71)	(23.85)	(24.1)	(27.33)	(26.39)
Female	0.00892***	0.00904***	0.00918***	0.00966***	0.00740***
	(15.94)	(16.15)	(16.4)	(18.46)	(16.41)
Female X Unemployment	0.000568***	0.000579***	0.000512***	0.000621***	0.000825***
	(6.00)	(6.12)	(5.41)	(7.00)	(10.62)
R^2	0.734	0.734	0.734	0.732	0.737
Observations	17,939,045	17,939,045	17,939,045	17,939,045	18,477,713

Notes: 1. This table presents the results for equation (1). 2. The dependent variable is a dummy indicating unemployment status, $y = \mathbb{1}(\text{non-working})$. 3. The first column displays results for the *female* and *race* variables in equation (1). 4. The second column includes results after adding state dummies to the control set ($\boldsymbol{x}$). 5. Columns 3, 4, and 5 present results with the addition of a time trend, state and time fixed effects, and 2-digit industry dummies, respectively. 6. The final column shows results using the state-level unemployment rate as the business cycle indicator (u). 7. Survey weights available in the CPS were used in the regression analysis. 8. T-statistics, calculated with robust standard errors, are shown in parentheses. 9. Key: *** significant at 1%; ** significant at 5%; * significant at 10%.

Table 10: Additional Robustness Checks Based on Age or Interactions.

		Sex		Age Range		Interactions and Age Range	
	Baseline	Male	Female	25 – 65	16 – 65	25 – 55	16 – 55
Black	0.0018*	-0.0043**	0.0069***	0.0028***	0.0107***	-0.0053***	0.0080***
	(2.06)	(-3.04)	(6.26)	(3.59)	(14.23)	(-3.77)	(6.85)
Black × Unemployment	0.0039***	0.0058***	0.0024***	0.0033***	0.0033***	0.0061***	0.0048***
	(25.97)	(23.94)	(12.90)	(24.76)	(26.29)	(25.38)	(24.45)
Female	0.0055***			0.0049***	0.0052***	0.0034***	0.0039***
	(10.52)			(10.51)	(11.77)	(6.08)	(8.25)
Female × Unemployment	0.00008			-0.00009	-0.0007***	0.0005***	-0.0004***
	(0.90)			(-1.17)	(-9.55)	(5.08)	(-4.36)
Black × Female						0.0132***	0.0051***
						(7.51)	(3.42)
Black × Female × Unemployment						-0.0040***	-0.0029***
						(-13.37)	(-11.39)
r2	0.041	0.046	0.038	0.040	0.041	0.042	0.041
N	18,567,096	8,943,998	9,623,098	22,335,305	27,901,175	18,567,096	27,901,175

Notes: 1. This table presents the results for additional robustness checks based on equation (1). 2. The dependent variable is a dummy indicating unemployment status, $y = \mathbb{1}(\text{unemployed})$. 3. The first column displays results for the $female$ and $race$ variables in equation (1) using a sample with age boundary specification of 25–55 (the baseline model). 4. The columns 2 and 3 show results for the model in column 1 but by sex. 5. Columns 4 and 5 present results after modifying in the model in column 1 the age boundary specifications. 6. The last two columns show results after considering additional interactions in equation (1) and varying the age boundary specifications in the sample. 7. Survey weights available in the CPS were used in the regression analysis. 8. T-statistics, calculated with robust standard errors, are shown in parentheses. 9. Key: *** significant at 1%; ** significant at 5%; * significant at 10%.