The Effects of Hiring Discrimination over the Business Cycle

Florian Kuhn* Luis Chancí†

August 19, 2021

The most recent version is available at http://www.floriankuhn.com

Abstract

Resume studies have found that certain demographic or social groups have lower callback rates for job interviews than others. In this paper we show that discrimination in hiring implies a higher volatility of labor market outcomes for the discriminated group in the context of a standard search-and-matching model with an urn-ball matching function. Intuitively, in recessions there are more applicants per job opening which hurts discriminated groups. In line with the model prediction, CPS data shows that black workers in the US have higher unemployment volatility over the business cycle compared to white workers when controlling for many observables visible to employers. We do not find the same effect for women when compared to men, consistent with the fact that resume studies generally find hiring discrimination for women to be at least an order of magnitude smaller than for blacks. Quantitatively, our theoretic setup allows us to directly use the point estimates from resume studies as parameter inputs for the differential in hiring rates in our model. Doing so, and calibrating to the US labor market, we find that the model can explain 70% of the extra business cycle volatility in the black unemployment rate.

JEL codes: E24, E32, J64, J71.
Keywords: Unemployment, Discrimination, Business Cycle.

*Department of Economics, Binghamton University. Email: fkuhn@binghamton.edu
†Escuela de Ingeniería Comercial, Faculty of Business and Economics, Universidad Santo Tomás, Chile. Email: luischanci@santotomas.cl

1 Introduction

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 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 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 is correlated with group status like, for example, education.

In our model we take the level of hiring discrimination as given. Staying as close as possible to the evidence from the quasi-experimental studies, as 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. To investigate the cyclical implications of such different hiring rates we modify an urn-ball matching function to allow for arbitrary degrees of hiring discrimination, and embed it in a search-and-matching model of the labor market. Over the business cycle, the model predicts that the discriminated group suffers from higher unemployment volatility; that is, when the economy enters a recessions the unemployment rate among discriminated workers increases more strongly. The intuition of 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 then turn to CPS data and examine the volatility of unemployment rate and job-finding rate for two particular demographic groups, women and blacks, since many resume studies investigate the degree of hiring discrimination for those groups. The goal is to see if these volatilities are larger than their counterparts for the groups of whites and males, respectively, and if so by how much. Importantly, we focus on *conditional* employment rates and job-finding rates; that is we calculate the respective probability of being employed and finding a job controlling for many observable characteristics. In line with previous literature we find strong differences in unemployment volatility for blacks compared to whites; on the other hand 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.

Finally, we calibrate the search-and-matching model with two types of workers to study the equilibrium effects of different hiring rates quantitatively. The way we set up the matching function allows us to use the difference in conditional hiring discrimination as a parameter, and 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 rates and job-finding rates, the model indicates that discrimination rates at a level found in resume studies could explain more than half of blacks' extra volatility in unemployment rates over the cycle.

Our focus on the hiring margin of course 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) and Fang and Moro (2011)). In this paper we 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.

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 receives 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 (we will discuss the difference between callback rates and hiring rates in more detail below). 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 works also 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 minority workers increases stronger than the unemployment rate for whites. The basic intuition of this is as follows. In a recession the labor market is slack 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.

In the empirical part, we use the Current Population Survey (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[1], 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 change of unemployment increases by 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 this 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.

The main contribution of this paper is to provide a 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.[2] Leveraging the structure of the standard search-and-matching model of the 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. We also contribute to a relatively thin literature of incorporating racial and gender heterogeneity into the structure of a model focusing on aggregate outcomes.

The outline of the paper is as follows: In the remainder of this introduction we review some of the most relevant work in a large body of research on discrimination. In section 2 we lay out the basic mechanism of hiring discrimination formalized by an urn-ball matching function, and

[1]Like immigration background, sexual orientation, parenthood, military status and many more, see for example Baert (2017).

[2]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.

incorporate it into a basic search model of the labor market, showing qualitatively how it leads to cyclical differences in labor market outcomes. Section 3 then analyzes labor market differences empirically using the CPS. In Section 4 we use the empirical findings to calibrate the model and assess quantitatively its implications for the labor market impacts of hiring discrimination over the business cycle. section 5 concludes.

Related literature One of the strands of literature this paper is related to is the search theoretic 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 which in turn endogenizes negative duration dependence of unemployment exit rates for an individual. In this paper we study 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.

Finally this paper 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 main findings of these types of experiments are surveyed in Bertrand and Duflo (2016), Neumark (2018), while Baert (2017) aims to collect all 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.[3] In section 4 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, 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

[3] 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).

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 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).[4] In the present paper we are correspondingly agnostic about the nature of discrimination.

Finally, 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[5]. 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. 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.

2 Model

To investigate the business cycle effects of discriminatory hiring formally, we develop a search-and-matching model with an urn-ball matching technology in the spirit of Blanchard and Diamond (1994). This matching mechanism at the heart of the model captures competition between workers and differential preferences by employers for different types of workers. We generalize Blanchard and Diamond (1994)'s setup to allow for arbitrary degrees of discrimination: in their model, whenever two workers of different groups compete for the same vacancy the worker from the preferred group *always* gets the job. In contrast, if the two workers are in the same applicant pool in our model, the worker from the preferred group has a higher chance of getting the job (but not necessarily an infinitely higher chance). We hence capture the degree of discrimination as the relative hiring probability between two candidates conditional on being in the same ap-

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

[5]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)

plicant pool, and we assign it to a key parameter in the model. Notably, this parameter has the same interpretation as the object of interest in resume studies, except that in those studies it is not the relative hiring probabilities that are directly observable but the relative probabilities for callbacks.

2.1 Environment

We briefly describe the economic environment most of which we keep standard, before turning in more detail to the matching function in the next subsection.

Time is discrete. There is a unit mass of workers, N_1 of which are in demographic group 1 and $N_2 = 1 - N_1$ of which are in group 2. Group membership is the only source of heterogeneity; in particular, there are no differences in productivity across workers. A worker of group i receives wage income w_i (we willl describe the wage-setting process in section 2.3), whereas unemployed workers receive an unemployment 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. All employed workers stay in their current job until they get separated from their match with an exogenous probability s. Workers consume all of their income in the current 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, so that in equilibrium the total number of active firms will be determined by a zero-profit condition. The cost of holding a vacancy open is c per period, and the probability of matching with a job-seeker is q_i which will again be determined by the matching technology as an equilibrium object. If the firm matches with and hires a worker, it will produce output y for every period that the match persists. The amount of output y produced is constant across firms, workers and matches, but does vary stochastically over time, generating business cycles.

2.2 Matching function

The matching technology determines the job-finding probabilities p_i and worker-finding probabilities q_i as a function 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 type of setup, every application by a worker is represented by a ball and every vacancy by an urn. Every period, in the application stage every unemployed worker submits one application to one of the posted job openings at random – figuratively, every ball gets randomly placed in one of the available urns. If there are many urns and balls, a law of large numbers guarantees that there is a fixed distribution of balls across urns; in other words there will be a certain fraction of urns with zero balls, a certain fraction of urns with exactly one ball, and so on. 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. On average, however, employers have a bias to hire from one of the two groups; say without loss of generality that they are biased to favor group-1 workers over workers from group 2. We fix the relative probability that a given worker from group 1 is picked for a job relative to a given worker from group 2 as the parameter π. For example, let $\pi = 2$, and consider an applicant pool that contains a given group-1 worker and a given group-2 worker: Then the group-1 worker's chances of getting the job are twice as high as the group-2 worker's, independent of the size and makeup of the remainder of the applicant pool. (Of course in the simplest case where these two are the only candidates in the pool this implies that respective hiring probabilities are 2/3 and 1/3.)

We take the probability π as given and constant over the business cycle. One easy way to interpret it is as the result of taste-based discrimination: We can think of employers making a logit-type choice among applicants, basing the decision on a latent, match-specific random variable that is the hiring manager's personal preference unrelated to the applicant's productivity. This latent preference variable may, however, be correlated with the applicant's group membership. We also 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.[6]

In the context of the urn-ball model the applications of different workers are represented by different types of balls, for example red and white. Formally, let Ω be the number of urns and Υ the number of balls, Υ_1 of type 1 which are red and Υ_2 of type 2 which are white. Every ball will be placed in an urn at random with uniform probability across urns. Define the ratio of balls to urns as the market tightness $\theta = \frac{\Omega}{\Upsilon}$, and $\theta_1 = \frac{\Omega}{\Upsilon_1}$ and $\theta_2 = \frac{\Omega}{\Upsilon_2}$ tightnesses with respect to each type of ball, respectively. Because all balls are placed independently from each other, the number of balls assigned to any given urn follows a binomial distribution. As such, if both Ω andΥ are large it can be approximated by a Poisson distribution with 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 colors of balls, because all balls are distributed indepen-

[6]The basic idea behind 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 iid signal about an applicant's quality. For two candidates with the same underlying productivity, a group-2 candidate needs a higher signal to be hired over a group-1 candidate resulting in differential hiring probabilities out of the same pool. Our conjecture is that those relative probabilities do not move strongly over the cycle. This will be the case as long as there is not too much movement in the group-specific means among unemployed workers.

dently the probability of having k_1 red balls and k_2 white balls is simply the product

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

noting that $1/\theta = 1/\theta_1 + 1/\theta_2$ from the definition of the market tightnesses. By 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 gets drawn at random from every urn. The two types of balls have different probability of getting drawn out of a given urn; without loss of generality say that the red balls are (weakly) more likely to be drawn. Let $\pi \geq 1$ be the relative probability 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 probability of drawing a red and white ball, respectively, is

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

Again using the law of large numbers, the total number of red and white balls drawn out of all urns combined is

$$\begin{aligned}\#reds &= \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \Pr\left(k_1, k_2; \theta_1, \theta_2\right) \Pr_{1|k_1,k_2} = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{\pi k_1}{\pi k_1 + k_2} \\ \#white &= \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \Pr\left(k_1, k_2; \theta_1, \theta_2\right) \Pr_{2|k_1,k_2} = \Omega \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{k_2}{\pi k_1 + k_2}.\end{aligned}$$

Finally, define as p_1 and p_2 the probability for any red and white ball, respectively, to be drawn out of some urn. These probabilities are then given by the number of total balls drawn relative to all balls of the same color:

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

$$p_2\left(\theta_1, \theta_2\right) = \#whites/\Upsilon_2 = \theta_2 \sum_{k_1=0}^{\infty} \sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1} \theta_2^{k_2} k_1! k_2!} \frac{k_2}{\pi k_1 + k_2}. \tag{2}$$

Returning from the urn-ball analogy to the economic model, $p_i(\theta_1, \theta_2)$ are the respective job-finding probabilities which households can calculate knowing the number of vacancies and the number of unemployed workers of both groups (i.e. both market tightnesses θ_1 and θ_2). As equations (1) and (2) show, these probabilities are composed of the likelihood of competing with k_1, k_2 other applicants for the same job times the likelihood of being picked from that applicant pool; and integrating 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

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

2.3 Wage setting and value functions

What is left to specify to close the model is the wage-setting rule. In principle we are free to pick any such 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 main calibration in section 4 we will use a standard Nash-bargaining rule. As alternatives we also consider constant wages as well as a no-discrimination policy under which wages are required to be equal across groups.

We now summarize the model by stating the implied value functions. Together, the market tightnesses θ_1 and θ_2 as well as exogenous productivity y describe the state of the economy, and match-finding probabilities and wages are functions of those state variables. A worker of type i receives wage income and either continues in the same job or at the fixed probability s becomes unemployed[7]. The associated value function is

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

Unemployed workers receive their unemployment benefit and have a chance to find work, and otherwise continue unemployed next period:

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

Firms' current period return when in a match is given by output produced minus wages paid, and their continuation value is the expectation of staying in the match versus exogenously separating from the worker. In a match with worker of type i, a firm's value function is

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

An unfilled posted vacancy costs a firm an amount c per period and gets filled with a worker i

[7] Chancí-Arango (2020) relaxes the assumption of constant separations rates to allow for endogenous separations and finds that the differences are small.

with probability $q_i(\theta)$. The value of a vacancy is hence

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

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

Appendix A collects all the equilibrium conditions.

2.4 Business cycle effects

In this section we analyze qualitative properties of the model to build intuition about the mechanism. We leave the full numerical evaluation and dicsussion of the calibration for section 4.

Impulse response functions Figure 1 displays the response of the model following a shock to output per match y via impulse response functions. Consider first the initial periods pre-shock when the economy is in 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 is reflected in higher steady-state job-finding rates for group 1 than for group 2, which in turn leads to 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. Because hiring exceeds separations for a few periods the unemployment rate decreases before recovering gradually in a U-shaped way.

Of main interest are the differences in unemployment rates and job-finding rates between the two groups (which are 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 group 2 unemployment decreases by more, reducing the difference between the two groups' unemployment rates.

Effect of individual parameters on labor market dynamics We can investigate how the different model parameters affect the differences between groups in response to a shock. Figure 2 display differences between the groups for different values of the parameter π, which determines the relative odds of getting hired for two applicants from different groups who are in the same applicant pool. The case that $\pi = 1$ means that workers of both types have the same chance to get hired out of a given pool. Hence there are no differences between the groups in the steady-state unemployment or job-finding rates; and both groups respond identically to a business cycle shock. For values of π strictly greater than 1, workers of group 1 get hired more readily, at the expense of group 2 applicants. As described above, the discriminated group's unemployment rate responds stronger 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 difference in the impulse response between groups.

Figure 1: Impulse responses to positive productivity shock

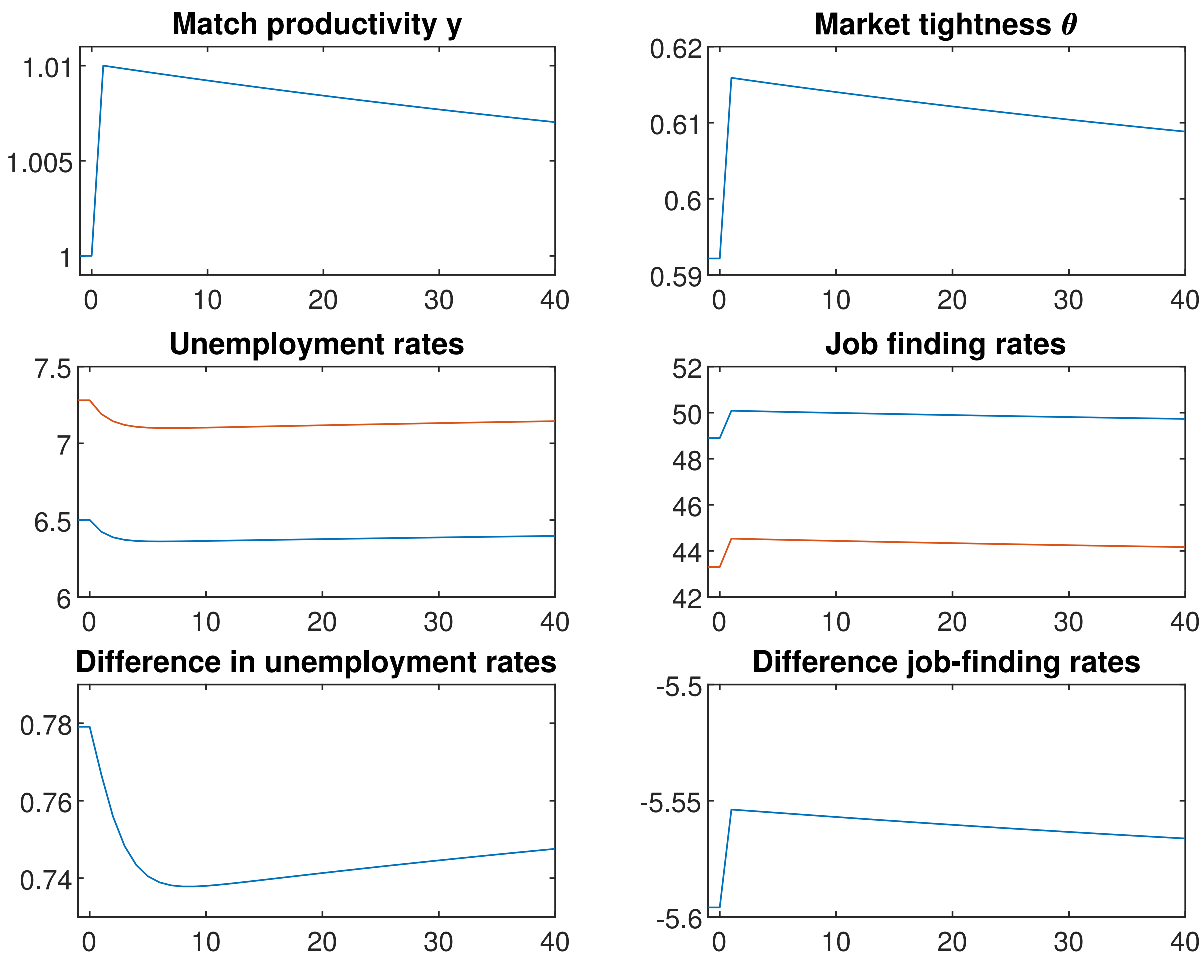

Notes: Periods of time in x-axis.

Note, however, that 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. With this extreme degree of discrimination the model nests Blanchard and Diamond (1994)'s case of lexicographic employer preferences.

We also consider variations in the cost of job creation c and the separation rate s (see 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. As a consequence unemployment is higher and job-finding rates are lower in steady state; but an increase in these parameters also implies higher volatility in market tightness for a given shock to y since it raises how much a firm benefits additionally from filling a vacancy. In the model, in turn, greater volatility of θ translates into a larger difference in the groups' response via the congestion mechanism.

3 Empirics

We use the Current Population Survey (CPS) to assess the relative volatility of different groups' unemployment rates over the business cycle. We focus on women and blacks for two main reasons: First, there are many resume studies investigating potential bias in the hiring process of these groups. Second, gender and race are characteristics that are straightforward to work with empirically: it is measured in the CPS (unlike e.g. sexual orientation), and group membership is relatively clearly defined, and for most – though not all – individuals it is binary and stable (unlike e.g. immigration background or disability). Finally, these characteristics are not endogenous to labor market conditions or employers' hiring rates themselves (like e.g. long-term unemployment or parenthood)[8]. The exercise here is similar to Cajner et al. (2017) and Hoynes et al. (2012).

The CPS is a rolling panel of housing units, which are surveyed according to a 4-8-4 pattern: Residents remain in the data set for four consecutive months, drop out for the following 8 months, and then are surveyed again for a final four months. Because housing units are the unit of observation, and not residents, directly one limitation of this dataset is a possible concern about sample selection (e.g. if a previously unemployed person moves because she found a new job that causes her to drop out of the sample), to the extent that it affects different groups differently. Within a household all persons are surveyed.

We use monthly data available from 1984 to 2018, and drop persons who are younger than

[8]This latter point is useful 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).

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

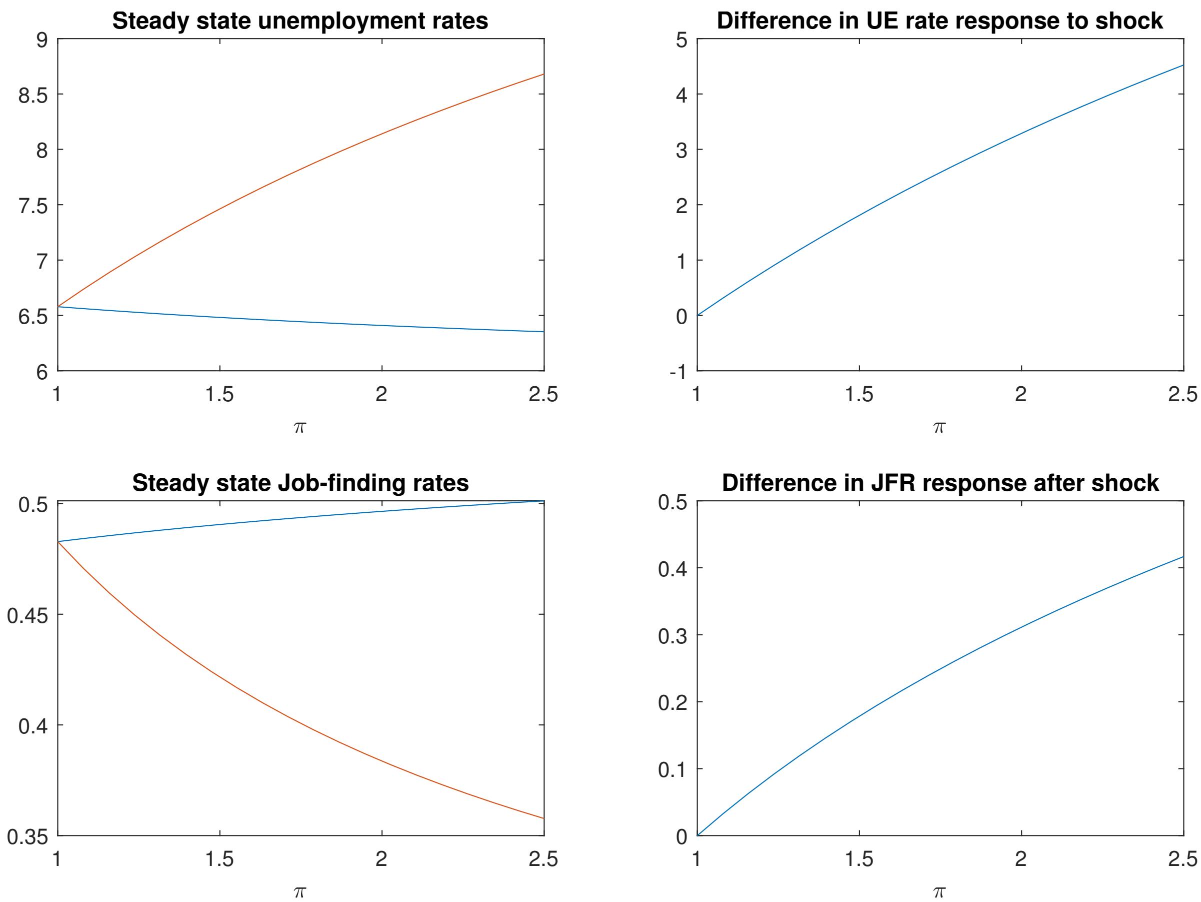

Notes: Model results for different values for the discrimination rate π (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity y in unemployment (top right) and job-finding rates (bottom right). 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. For the full calibration and numerical results, see section 4.

25 or older than 55 years old, who are retired or who are members of the armed forces. For our baseline analysis we run pooled OLS, using as outcome variable an unemployment dummy which takes the value of 1 if the individual does not have, but is looking for, a job. As a measure of the state of the business cycle we use the aggregate unemployment rate as provided by the BLS. In the regression we will focus on the coefficients on the interaction of the sex/race dummy with the unemployment rate. These coefficients indicate how much the group-specific likelihood of employment (conditional on the demographic variables described below) increases relative to a white male's likelihood of unemployment when the aggregate unemployment rate increases by one percentage point.

While with observational data we will not be able to exactly mimic the setup of resume studies which assign gender or race randomly and hold all other factors constant, we control for a number of demographic characteristics that are captured in the CPS. Importantly, these variables are generally also observable by employers during the hiring process so that we do not condition on something that employers can not. These regressors are (a quadratic in) age, educational achievement[9], family status[10], city size and occupation in the detailed categories provided by the CPS. The latter variable refers to the occupation in which the individual is currently employed or, in case of unemployment, held last. 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.

Equation (3) summarizes our main regression specification for a person i at time t:

$$Y_{it} = \boldsymbol{X}_{it}\boldsymbol{\beta}_1 + U_t\boldsymbol{X}_{it}\boldsymbol{\beta}_2 + \gamma_1 \text{Black}_{it} + \gamma_2 \text{Black}_{it} U_t + \gamma_3 \text{Female}_{it} + \gamma_4 \text{Female}_{it} U_t + \varepsilon_{it} \tag{3}$$

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.

Table 1 displays the estimation results for the coefficients on race and gender, respectively, i.e. γ_1 through γ_4 of equation (3). Column 1 contains the main specification just described, and our interest is on the interactions of group status with the aggregate unemployment rate. The coefficient on "Black $\times$ Unemployment" of 0.00394 means that on average for every percentage point increase in aggregate unemployment, the unemployment rate of blacks increases about 0.4 percentage points more than the one of whites. In other words, assume that in a severe recession aggregate unemployment increases from 5 to 10 percent. Using average values for the other control variables, the gap of unemployment for blacks with respect to whites will sharply increase from 2 to 4 percent, which is an economically large effect. On the other hand we do not find statistically nor economically large effects for women. Viewed through the lens of our

[9] Less than high school, high school, some college, college, post-graduate degree.

[10] Married without children, married with children, unmarried without children, unmarried with children.

Table 1: Unemployment Status and Business Cycle.

	Baseline	State Dummies	Time-trend	State and Time FE	Industry Dummies	State Unempl.
Black	0.00162	0.00272**	0.00157	0.00269**	0.00433***	0.00279***
	(1.77)	(2.92)	(1.72)	(2.89)	(4.62)	(3.63)
Black × Unemployment	0.00394***	0.00404***	0.00394***	0.00404***	0.00464***	0.00363***
	(25.49)	(25.6)	(25.52)	(25.55)	(29.27)	(28.04)
Female	0.00550***	0.00560***	0.00552***	0.00563***	0.00579***	0.00416***
	(10.18)	(10.37)	(10.21)	(10.43)	(11.48)	(9.32)
Female × Unemployment	0.0000746	0.0000857	0.0000748	0.0000873	0.000102	0.000321***
	(0.81)	(0.93)	(0.81)	(0.95)	(1.19)	(4.19)
R^2	0.0415	0.0426	0.0415	0.0428	0.0365	0.0439
Observations	17,939,045	17,939,045	17,939,045	17,939,045	17,939,045	17,939,045

Notes: The table reports the results for equation (3). The dependent variable is a dummy variable for unemployment status. First column shows the results for $Female$ and $Race$ in equation (3). Second column contains the results after adding state dummies to the set of controls ($\boldsymbol{X}$). Columns 3, 4, and 5 are the results after adding a time trend, state and time fixed effects, and Industry 2-digits dummies, respectively. Last column presents the results after using state-level unemployment rate as indicator for business cycle (u). t-statistics in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

model, this is consistent with the fact that resume studies find large differences in callback rates for blacks, but only small (if any) differences in callback rates for women.

This overall pattern extends throughout a set of robustness checks, listed in the remaining columns of table 1. These include the addition of state dummies, alternative industry controls, or the use of state-level unemployment rate instead of the aggregate unemployment rate. Since it is well established that labor force participation has strong cyclical components, as an additional robustness check we also consider using the non-employment rate (individuals either unemployed or not in the labor force) instead of the unemployment rate as the outcome measure (see Appendix D). Overall we find a very similar pattern.

Table 2: Job-finding Rate and Business Cycle

	Job-finding Rate
Black	-0.0637***
	(-10.67)
Black × Unemployment	0.00459***
	(5.32)
Female	-0.0314***
	(-6.08)
Female × Unemployment	0.000554
	(0.74)
Observations	612,034

Notes: The table reports the results for job-finding rates defined as the rate of those who switch their unemployment status. t statistics in parentheses. Key: *** significant at 1%; ** 5%; * 10%.

In our model the mechanism through which differential unemployment rates come about is because of differences in hiring rates. We therefore also consider the behavior of job-finding rates over the cycle in table 2. The table shows that the gap between blacks and whites, in terms of the probability of finding a job out of unemployment does *not* significantly widen in a recession (the coefficient on the interaction is even positive, meaning the gap decreases as aggregate unemployment increases). However, there is a large difference in baseline *levels* of job-finding rates as evidenced by the large negative coefficient on the Black dummy. This, again, is 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.

4 Results

4.1 Calibration and results

We calibrate the model to the US economy using *aggregate* labor market statistics; that is, not taking any group differences into account. Specifically, we target the long-run average level of the (aggregate) unemployment rate and the average job finding rate. We then use the model to investigate the cyclical differences in labor market outcomes between groups (in this case the differences between blacks and whites) for a given level of hiring discrimination π. In other words, we look at how the groups' unemployment rate and job-finding rates respond following aggregate shocks, and how large a difference in these responses a given value of π generates.

Table 3 collects the parameter values used. Given the data's monthly frequency we choose $\beta = 0.9967$ which corresponds to an annual discount factor of 0.96. Following Shimer (2005a),

Table 3: Calibration

Parameter	Value	Source
β discount rate	0.9967	Monthly frequency, annual interest rate of 4%
s separation rate	0.034	Average separation rate (Shimer, 2005a)
b value of unemployment	0.71	Standard value, e.g. Hall and Milgrom (2008)
ρ_y persistence of productivity	0.983	Quarterly autocorrelation of output 0.95
σ_ε sd 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	0.46	Calibrated to mean agg UE and JF rate
π degree of hiring discrimination	1.38	Resume studies

Notes: Parameter values used in the model. 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, 2005a).

we use a separation rate of $s = 3.4\%$, and we set the flow value of unemployment to $b = 0.71$ following Hall and Milgrom (2008), which is also in the middle of the range of common estimates for this parameter. In the CPS the ratio of blacks to whites among labor force participants is roughly 1 to 9.

Finally, we calibrate both the vacancy creation cost c and the employer's bargaining power ν targeting the long-run mean of the unemployment rate in the CPS of 6.2% and the job-finding rate as reported by Shimer (2005b) 45%. The implied value for the vacancy creation cost is to 0.46, a standard value in the literature corresponding to a vacancy costing roughly 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 the model the calibrated value for c implies that aggregate vacancy creation costs are 1.9% of aggregate output. For the employer's bargaining power the calibration yields $\nu = 0.564$. Again within the range of values commonly used in the literature, this implies that employers receive around a bit more than half of the joint match surplus. Table 4 shows that the model moments are reasonably close to the targets.

Finally, we pick the degree of hiring discrimination to be $\pi = 1.385$ for our baseline calibration, which by our count is the median estimate of the resume studies surveyed in Baert (2017)[11] focusing on African-Americans in the US. This value is also close to the point estimate of 1.49 in the seminal study of Bertrand and Mullainathan (2004). It is noteworthy that we can pick π independently from the other model parameters because it affects the aggregate behavior of the model only minimally. Instead its first-order effect is only on the *distribution* of labor market

[11]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.

Table 4: Model results

Outcome	Aggregate	Group 1	Group 2
Mean unemployment	6.58 %	6.50 %	7.28 %
Std dev unemployment	0.21 %	0.21 %	0.27 %
Mean job-finding rate	48.3 %	48.9 %	43.3 %
Std dev job-finding rate	1.69 %	1.69 %	1.75 %
Wages	0.9792	0.9794	0.9771
Std dev wages	0.2145	0.0136	0.0137

Notes: Steady state values and standard deviations of outcomes in the model.

outcomes between groups.[12]

Table 4 displays the model's results in terms of aggregate outcomes, and outcomes by groups. Most notably, table 4 shows that there are both strong level effects and strong cyclical effects of hiring discrimination on unemployment. With an unemployment rate of 7.5%, group 2 has a 0.85 percentage points higher unemployment rate than group 1. Similarly, it takes workers of group 2 significantly longer to find a job if they are unemployed, with a substantial gap in job-finding rates of 5.8 percentage points. The cyclical movements in the unemployment rate, which this paper focuses on, are much stronger for group 2. Group 2's unemployment rate has a standard deviation of 0.925 percentage points, whereas the standard deviation of group 1's unemployment is 0.705 – in other words, fluctuations in group 2's unemployment rate exceed group 1's by 31%. It is worth pointing out that these adverse cyclical effects for group 2 do not show up in a significantly higher volatility of job-finding rates. This is because 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.

To put the numbers of table 4 into context and compare them to our empirical findings of section 3, 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 difference in unemployment rates increases by 1.39 percentage points, thus accounting for about 70% of our empirically measured gap.

Finally, hiring discrimination does not seem to explain the racial *wage* gap quantitatively.

[12] 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.

It is true that qualitatively the mechanism generates a difference in wages between groups, and 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 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) put this gap at around 9% for men and 5% for women.

4.2 Counterfactuals

We can use the mapping from hiring discrimination to the difference in labor market outcomes that the model provides to answer counterfactual questions. Two questions of interest are, how large is the degree of hiring discrimination which we would have to assume to explain the full difference in labor market outcomes? And second, what is the effect of reducing discrimination? In other words, by how much does the gap in labor market outcomes narrow if hiring discrimination is reduced?

To answer the first question, we increase the value of π until the relative difference in unemployment volatilities corresponds to our empirically measured value – this is the case at a value of $\pi = 1.57$. This is a value close to Bertrand and Mullainathan (2004)'s estimate of 1.49 which in turn is still well within the range of the other resume studies surveyed in Baert (2017). We can think of this exercise as trying to identify the degree of hiring discrimination off of 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 for a steady-state difference in job-finding rates of 7.6 percentage points (compared to 5.8 points before).

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 0.33 percentage points from 7.28% to 6.95%, and an increase in steady-state job-finding rates of 2.2 percentage points from 43.3% to 45.5%. Unemployment volatility is reduced and now exceeds the one of whites by 15.2%, instead of 29.7% in the baseline calibration[13]. This reduction in volatility means that in the case of our exemplary big recession with 5% higher aggregate unemployment, the black unemployment rate increases 0.73 percentage points stronger than for whites', compared to 1.39 percentage points in the baseline. In other words, the reduction in hiring discrimination could shave off 0.66 percentage points of black unemployment in a severe recession.

[13] As discussed in section 2 the difference in labor market outcomes is non-linear in π, but this concavity is not very pronounced in this area of the parameter value.

4.3 Alternative wage-setting rules

We consider two plausible alternatives for wage determination in the model to investigate to which extent wage setting affects the results. The first is a no-discrimination rule under which employers are not allowed to condition wages based on group status, and the second is constant wages in which wages do not move with productivity. In other words, while our baseline Nash bargaining rule conditions wages on group status and business cycle conditions, alternative 1 removes variability of wages between groups and alternative 2 removes variability of wages over the cycle, allowing us to assess the impact of these implications of the Nash rule separately.

No wage discrimination For the first alternative we assume that wages at any time have to be equal across groups. This means that at any point of the business cycle there is only one wage which can depend on productivity, but not on group membership. Of course this is not consistent with standard Nash bargaining because the two worker groups have different respective outside options (recall that group 1's value of unemployment is higher due to the shorter expected unemployment duration). We therefore model the wage as sharing the joint surplus between an employer and the *average* worker:

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

To the employer there is now no difference of being matched so we can drop the index i on the value of a match for the employer J.

This wage rule, because it has minimal effects on the *aggregate* wage level (i.e. the cost of postig a vacancy for employers), has barely any employment effects. In this setup, the equal-wage requirement has purely redistributive consequences: wages of group 1 decrease slightly to the benefit of group-2 wages, and there are no meaningful employment effects for either group. This rule therefore has the potential to undo the (small) negative wage effects of hiring discrimination. Of course, the direct effect of hiring discrimination on the differential employment dynamics persists.

Constant wages Under this setup, wages are constant over time and across workers of 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 $\overline{w}_2$.[14] It is well-known (e.g. Hall, 2005) that rigid wages can amplify the volatility of other labor market outcomes, so 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).

[14] 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.

Table 5: Alternative wage-setting rules

Outcome	Baseline (Nash) Group 1; Group2	No-discrimination wage	Constant wages Group 1; Group 2
Mean wages	.9794; .9771	.9792	.9794; .9771
Std dev wages	.0136; .0137	.0136	0
Stddev unemployment	.21; .27	.21; .27	1.13; 1.46
Add'l group-2 UE in rec	1.39	1.39	1.35

Notes: Comparison of the three alternative rules determining the wage. 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?

The results in table 5 show that as expected eliminating wage adjustment over the business cycle increases the volatility of unemployment fluctuations. However it does so fairly evenly for both groups, so that their relative standard deviations remain approximately the same as under Nash wages.

5 Conclusion

We extend the urn-ball matching function to allow for an arbitrary degree of hiring discrimination. Incorporating this matching function into a search-and-matching model implies higher unemployment volatility for the discriminated group. Using US data, and in line with previous research, we do find high unemployment volatility for black labor market participants relative to whites.

At the same time the model provides a quantitative mapping from the degree of hiring discrimination into differences in labor market outcomes. Using resume studies to gauge the existing degree of hiring discrimination in the US labor market numerically, the model generates around 70% of blacks' excess unemployment volatility. This suggests that, in addition to the well known effects on the level of unemployment, hiring discrimination has potentially large adverse effects on the business cycle behavior of unemployment rates.

References

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

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.

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.

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

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.

Chancí-Arango, L. D. (2020). Essays on Applied Econometrics. *Dissertation, Binghamton University*.

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.

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

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.

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.

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.

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.

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.

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 Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review 95*(1), 25–49.

Shimer, R. (2005b). The cyclicality of hires, separations, and job-to-job transitions. *REVIEW-FEDERAL RESERVE BANK OF SAINT ...*, 493–508.

Appendix A Equilibrium conditions

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

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

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

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

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

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

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:

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

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

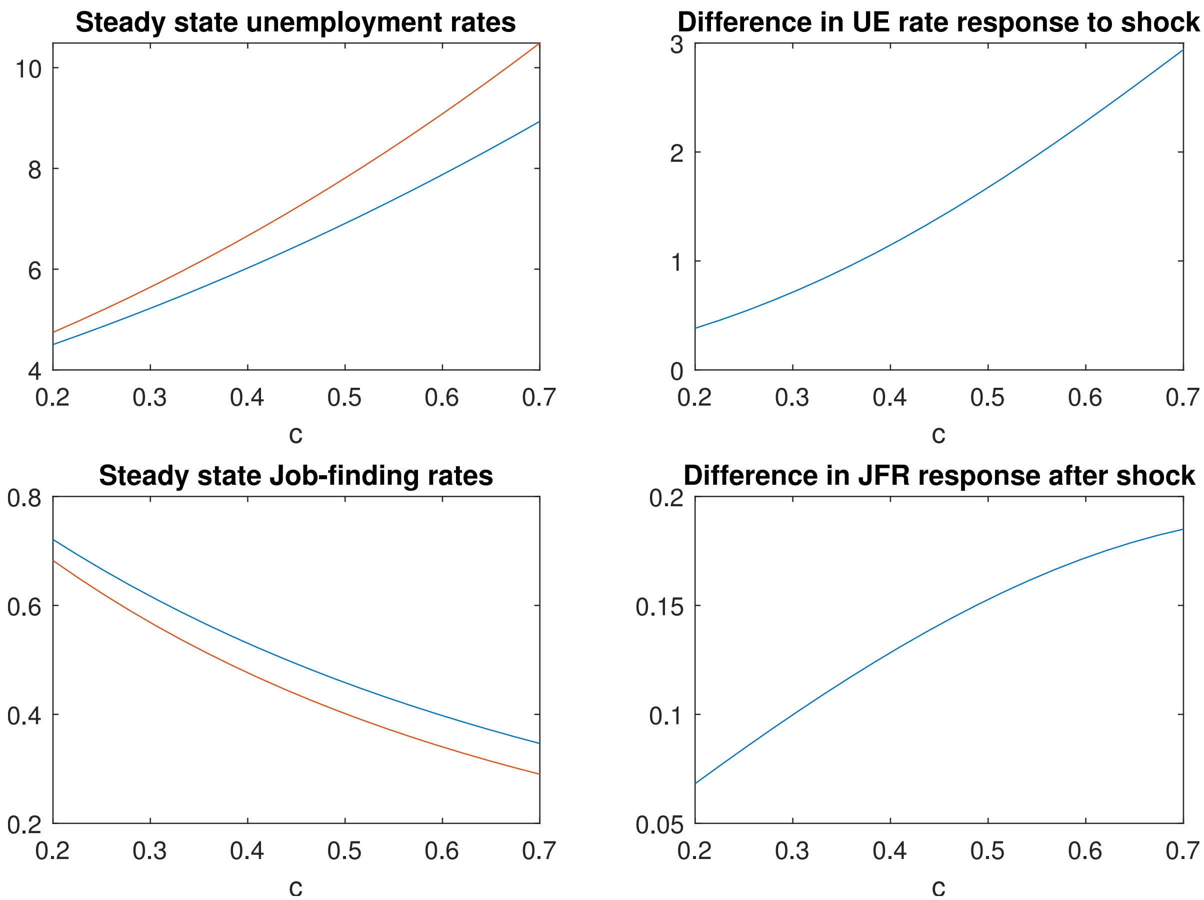

Notes: Model results for different values for the vacancy creation cost c (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after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

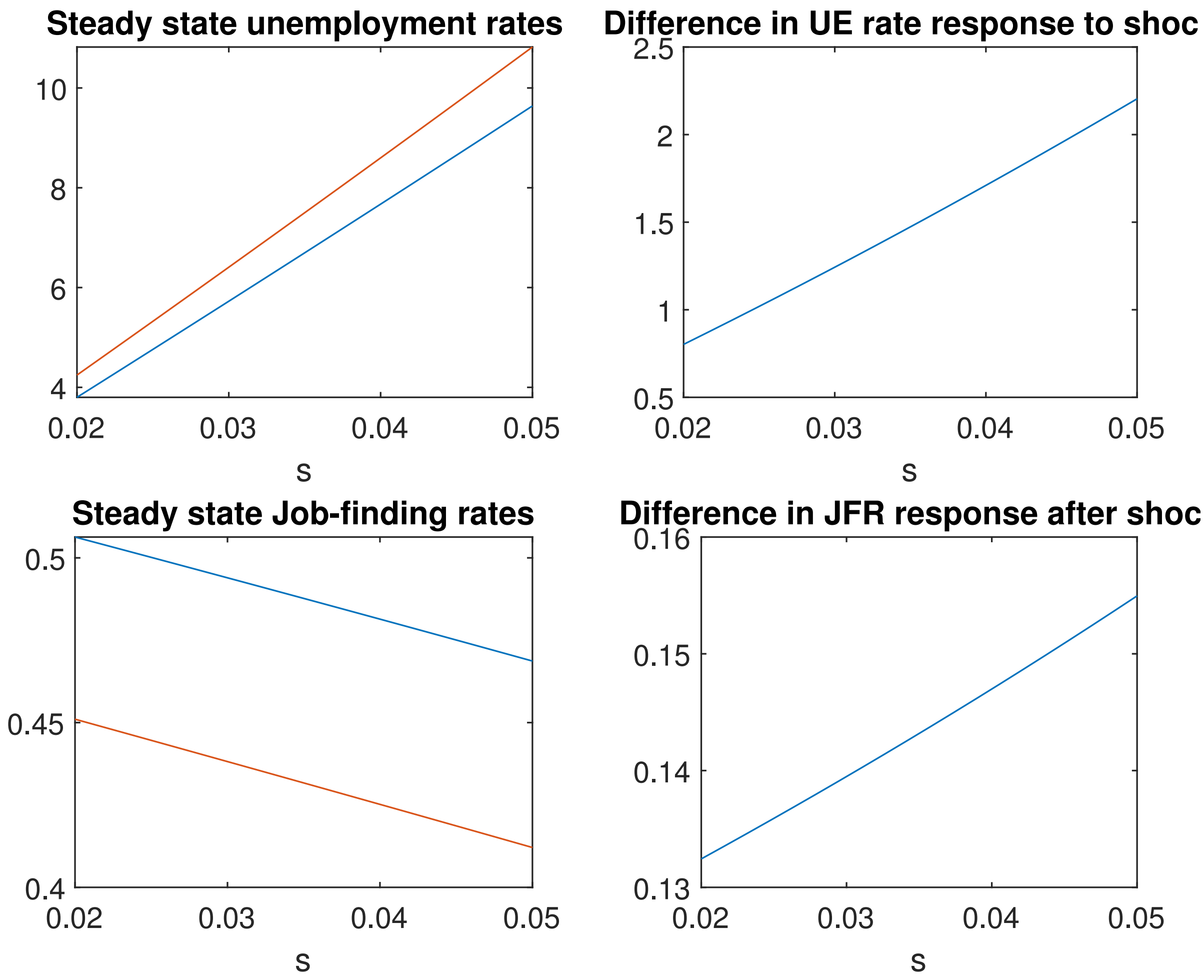

Notes: Model results for different values for the separation rate s (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase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

Appendix C Descriptive Statistics

Table 6: DESCRIPTIVE STATISTICS. 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: This table reports the mean value and the standard deviation (in parentheses) for the unemployment rate (in percentage, %) over 1984m1-2018m3. Author's calculations using CPS.

Appendix D Additional Robustness Checks

Table 7: Non-working Status and Business Cycle.

	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: The table reports the results for equation (3). The dependent variable is a dummy variable for non-working status. First column shows the results for $Female$ and $Race$ in equation (3). Second column contains the results after adding state dummies to the set of controls ($\boldsymbol{X}$). Columns 3 and 4 are the results after adding a time trend and Industry 2-digits dummies, respectively. Last column presents the results after using state-level unemployment rate as indicator for business cycle (U). t-statistics in parentheses. Key: *** significant at 1%; ** 5%; * 10%.