Effects of Hiring Discrimination over the Business Cycle

Florian Kuhn[*] Luis Chancí[†]

November 15, 2018

Comments are welcome

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. We show theoretically that such a form of hiring discrimination 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. In line with this prediction, we find in CPS data that blacks in the US have higher unemployment volatility over the business cycle compared to whites 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 an order of magnitude smaller than for blacks. Quantitatively, our theoretic setup allows us to directly use the point estimates from resume studies for the differential in hiring rates as parameter inputs in our model. Doing so, and calibrating to the US labor market we find that the model can explain most 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
[†]Email: lchanci1@binghamton.edu

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 these studies 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 we modify a 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 discriminated group suffers from higher unemployment volatility; that is, when the economy enters a recessions the unemployment rate among discriminated workers increases more strongly.

We then turn to CPS data and examine the volatility of employment 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. We find strong differences in volatility for blacks, but weak or no evidence of extra volatility for women. This is 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 employ a search-and-matching model with two types of workers to study the equilibrium effects of different hiring rates. The matching function we use allows us to use estimates from resume studies directly as a parameter. The model thus provides a mapping between the degree of discrimination and the volatility of labor market outcomes, allowing us to assess the order of magnitude of the effect that hiring discrimination has on cyclical labor market outcomes. Calibrated to match *mean* differences in employment rates and job-finding rates between blacks and whites, the model indicates that discrimination rates as found in resume studies could explain more than half of blacks' extra *volatility* in employment rates and job-finding rates.

Throughout our theoretical analysis of the business cycle effects we hold fixed the degree of discrimination 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 that the white person receives the job – this 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).

An obvious direct implication of a lower likelihood of getting hired is a level effect implying higher average unemployment and a lower job-finding rate among women. But in addition there are also business-cycle effects: In recessions, the unemployment rate among blacks 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 worse for blacks than for whites: The effects of discrimination play out exactly when blacks compete directly with whites, and recessions are times of increased competition among workers leading to larger applicant pools and higher odds that a black applicant will compete against a white applicant for a given job. The result is a bigger drop in employment and in the job-finding rate for blacks during recessions. We model this effect formally by extending Blanchard/Diamond (1993)'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], empirically these are the cleanest to work with 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 employment rates and job-finding rates exhibit 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. These findings are consistent with the results of resume studies which tend to show that hiring discrimination is an order of magnitude larger for blacks than for women.

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[2] the portion of labor market outcomes that are due to discrimination is not. Leveraging the structure of the standard search-and-matching model of the labor

[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.

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 exactly 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: First, we extend Blanchard-Diamond's urn-ball model in a tractable way to allow for an arbitrary degree of hiring discrimination. Second, we contribute to a 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 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. Finally, 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 explaining racial and gender discrimination, in particular. These papers 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 and job-finding rates by race and gender.

Finally, this paper is also related to a large body of empirical literature on discrimination. Similar to that literature we do not theorize about employer motives for discrimination, but examine its consequences for the labor market. As mentioned, an large and important part of that literature are resume and audit studies, which are surveyed in Bertrand and Duflo (2016), Neumark (2016) and Baert (2017). 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. But empirical evidence for discrimination is not limited to resume studies: One example is Black and Strahan (2001) who exploit the natural experiment of banking deregulation in the 1970s and use a triple difference to show that increased competition lead to convergence of male and female wages, consistent with taste-based discrimination.

There are 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 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. The degree of discrimination could hence be stronger or weaker than the effects measured by resume studies. We still think that the effect size measures in resume studies is informative about the degree of discrimination for a given group. 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).[3] In the present paper we are correspondingly agnostic about the nature of discrimination. This will not matter if the degree of discrimination does not vary over the cycle.

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

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 this relative hiring probability between two candidates conditional on being in the same applicant 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 Matching function

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. Time is discrete. 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 randomly drawing a ball out of the respective urn.

We abstract away from any differences between workers other than group status. In particular there are no differences in worker productivity.

We therefore assume that an employer will pick between applicants of the same group with equal probability. On average, however, employers have a preference 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 Jack (a group-1 worker) and John (a group-2 worker): Jack's chances of getting the job are twice as high as John's, independent of the size and makeup of the remaining applicant pool. (Of course in the simplest case where Jack and John are the only candidates this implies that respective hiring probabilities are 2/3 and 1/3.) One way to interpret this model of discriminatory hiring is the following: We can think of employers making a logit-type choice among applicants, basing the decision on a latent, match-specific random variable[4], and all employers assigning a negative shift in that latent variable to candidates of group 2.

In the context of the urn-ball model one can think of different types of balls, one of which is larger (and hence more likely to be drawn out of an urn) than the other. 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 independently 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 because they are of different sizes. Without loss of generality say that the red balls are (weakly) larger. 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,

[4]This could be, for example, the hiring manager's utility, or a signal of expected match quality if employers can't observe applicant quality perfectly in a statistical discrimination setting

respectively, is

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

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

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

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:

$$
\begin{aligned}
p_1(\theta_1,\theta_2) = \#reds/\Upsilon_1 \quad &= \theta_1 \sum_{k_1=0}^{\infty}\sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1}\theta_2^{k_2}k_1!k_2!}\frac{\pi k_1}{\pi k_1 + k_2} \\
p_2(\theta_1,\theta_2) = \#whites/\Upsilon_2 &= \theta_2 \sum_{k_1=0}^{\infty}\sum_{k_2=0}^{\infty} \frac{e^{-\frac{1}{\theta}}}{\theta_1^{k_1}\theta_2^{k_2}k_1!k_2!}\frac{\pi k_2}{\pi k_1 + k_2}.
\end{aligned}
$$

These will be the respective job-finding probabilities which agents can calculate knowing the number of vacancies and the number of unemployed workers of both groups (and hence both market tightnesses θ_1 and θ_2).

2.2 Remaining model setup

The remainder of the assumptions follow a standard search-and-matching model. A worker of type i has the value function

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

The only way for a job match to end is through a separation for exogenous reasons which happens with constant probability s. Workers receive a wage w_i which will depend on market tightnesses as well as on the aggregate state of the economy which is captured by the random variable y, so that together θ_1, θ_2 and y describe the state of the economy. Workers have a discount factor of β and are risk-neutral.

When unemployed, workers receive an unemployment benefit of b and look for a job. As derived in the previous section, their job-finding rate is given by p_i and depends on both market tightnesses. An unemployed worker's situation is captured by the value function

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

When a job is filled a firm produces output of y goods per period. Output y fluctuates randomly over time, representing (aggregate) business cycle uncertainty (we abstract from worker-specific, firm-specific or match-specific heterogeneity). The firm's flow profit from employing a worker is hence $y - w_i(\theta_1, \theta_2, y)$, so that its value of a match with a worker of type i is

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

If the job match is destroyed, the continuation value for the firm is the value of an unfilled vacancy $V(\theta_1, \theta_2, y)$. An unfilled posted vacancy costs a firm an amount c per period and gets filled with a worker i with probability $q_i(\theta)$. Since all employers with at least one job applicant will match with a worker, the function

$$q(\theta) = 1 - e^{-1/\theta}$$

is simply one minus the probability that no workers get matched to the employer. The value of a vacancy is hence

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

As is standard we will assume that there is free entry by firms to post vacancies such that $V(\theta_1, \theta_2, y) = 0$ at all times.

What is left to specify is the wage-setting rule. In principle we are free to pick any such rule that shares the joint match surplus. In the next section, to illustrate the effects qualitatively we will consider the case that employers have all bargaining power. In the model calibration in section 4 we will consider the more general Nash-bargaining rule, and, as alternatives, constant wages as well as wages under a no-wage-discrimination rule. Appendix A collects all the equilibrium conditions.

2.3 Business cycle effects

In this section we analyze a simple version of the model to ex/plore its properties qualitatively. To simplify the analysis we start with a special case of the model in which employers have all the bargaining power and are able to extract the entire joint match surplus. In that case the wages for all workers and states are identically $w_i(\theta_1, \theta_2, y) = b$. For the firm this also means that the value of having a filled vacancy does not depend on the type of worker so that we can write $J_1 = J_2 \equiv J$ and $q_1 + q_2 = q$. The more general model of section 4 with a Nash wage bargaining rule nests this example.

Impulse response functions Figure 1 displays the basic qualitative properties 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 men than for women, which in turn leads to relatively higher steady-state unemployment rate among women.

As is standard, an increase in output per match causes firms to post more vacancies resulting in a jump up 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. 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.

Figure 1: Impulse responses to positive productivity shock

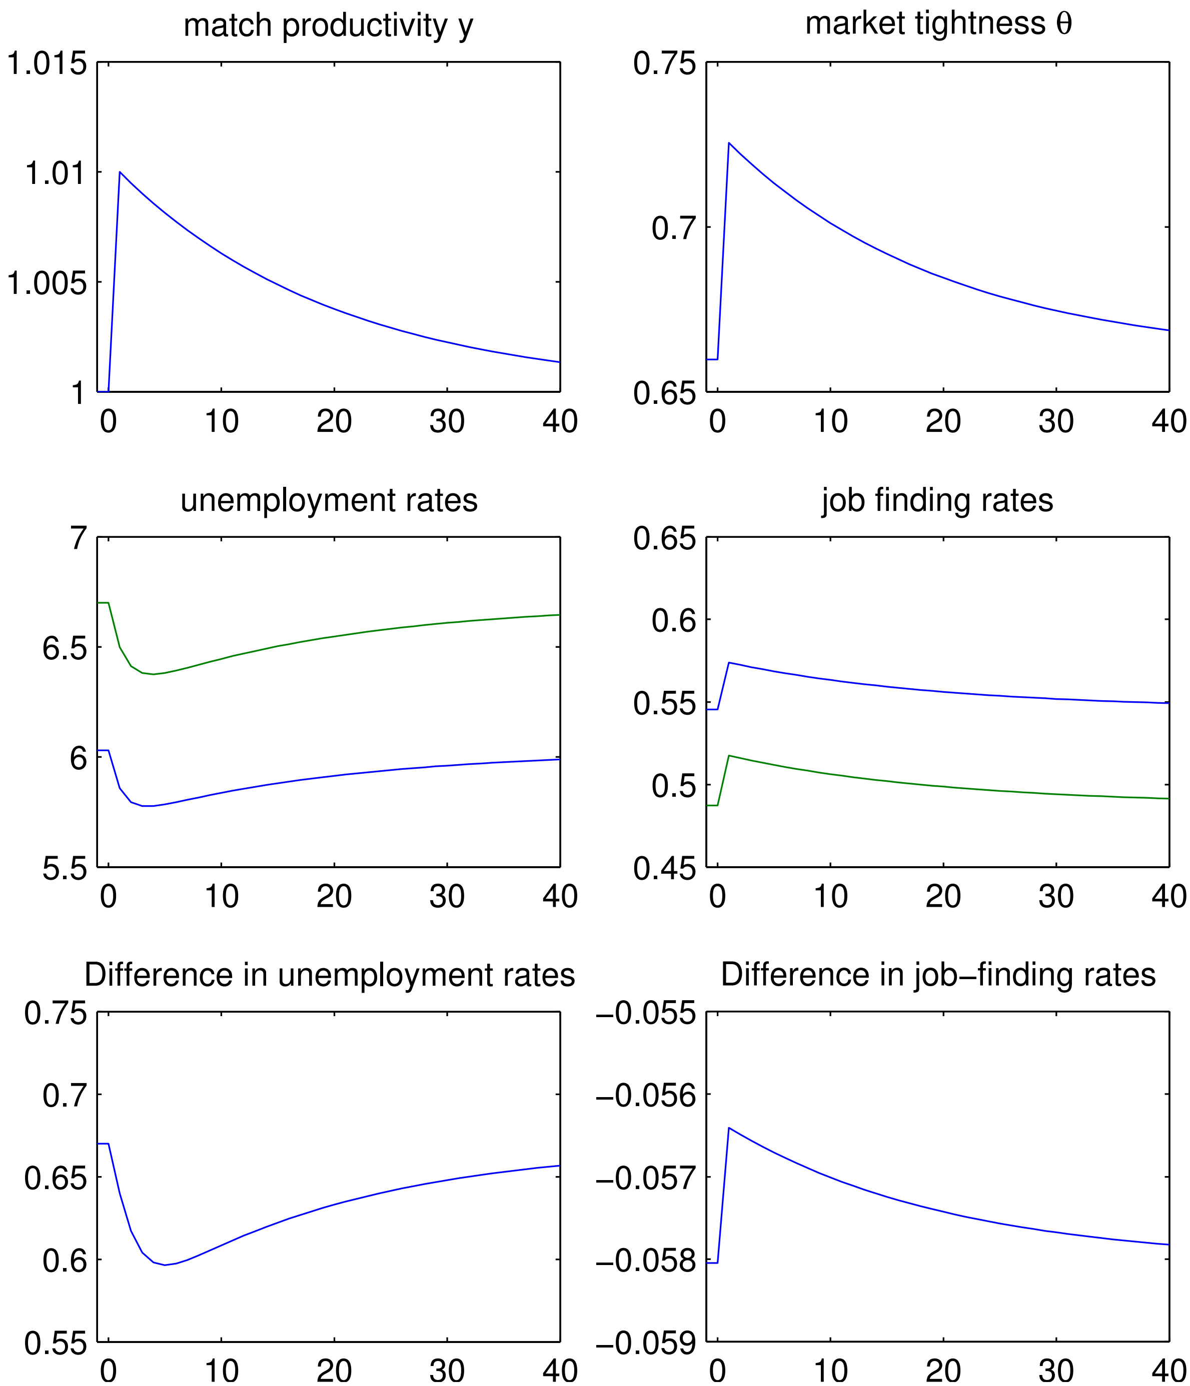

Notes: Periods of time in x-axis.

Figure 2: Group differences as function of discrimination π

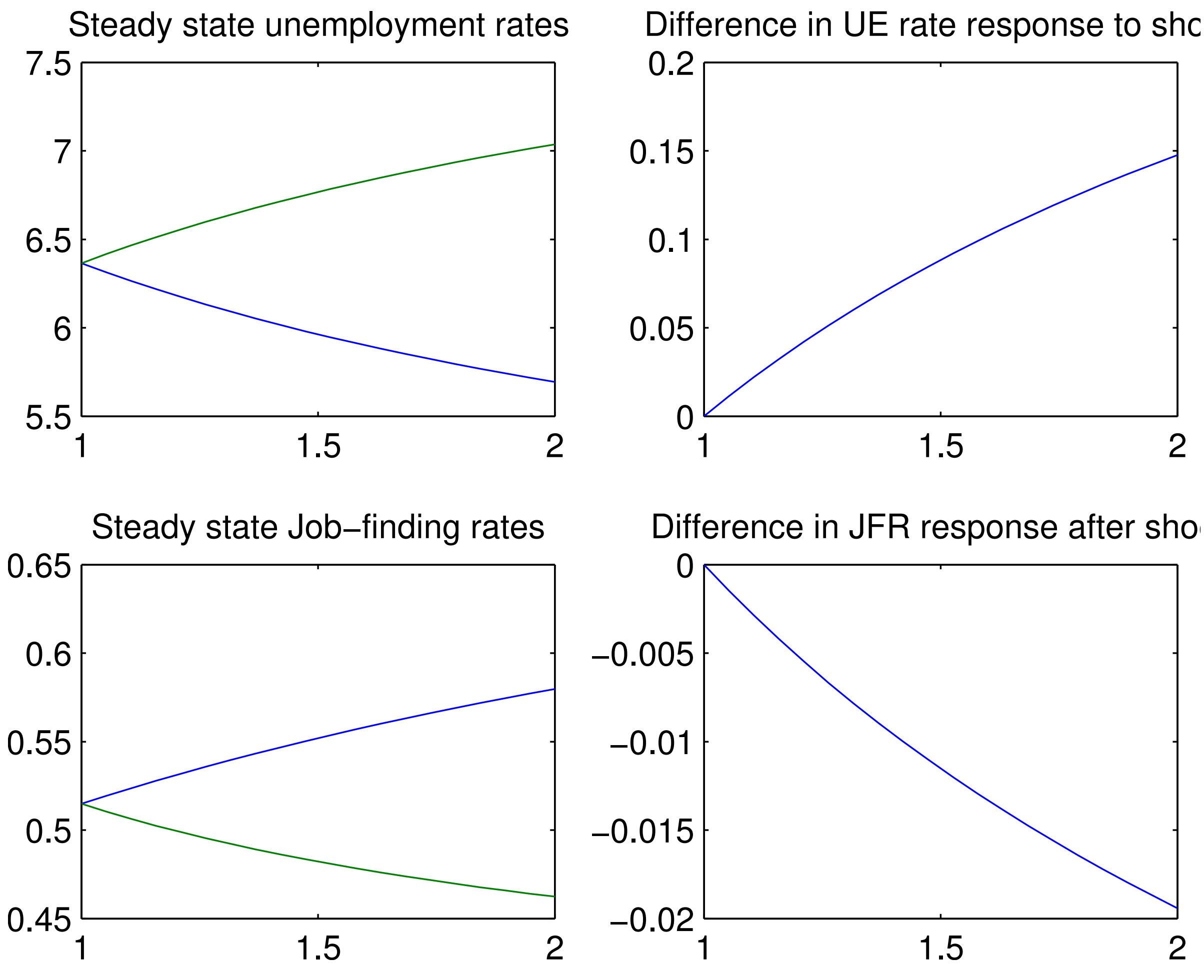

Notes: Model results for different values for π (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity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

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), group membership is well defined and binary (unlike e.g. immigration background or disability), and is not endogenous to labor market conditions or employers' hiring rates themselves (like e.g. long-term unemployment or parenthood)[5]. 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 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,

[5] 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, 2016, provide evidence that in recessions employers discriminate against long-term unemployed to a lesser degree, as is consistent with statistical discrimination).

educational achievement[6], family status[7], 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 genderrace with the state of the business cycle, all controls are included both in levels and as interaction with the aggregate unemployment rate.

Equation (1) summarizes our main regression specification for a person i at time t:

$$Y_{it} = \boldsymbol{X}_{it}\boldsymbol{\beta}_1 + U_t\boldsymbol{X}_{it}\boldsymbol{\beta}_2 + \gamma_1\text{Black}_{it} + \gamma_2\text{Black}_{it}U_t + \gamma_3\text{Female}_{it} + \gamma_4\text{Female}_{it}U_t + \varepsilon_{it} \qquad (1)$$

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.

Table 1: Unemployment Status and Business Cycle.

	Baseline	State Dummies	Time-trend	Industry Dummies	State Unempl.
Black	0.00162	0.00272**	0.00157	0.00433***	0.00274***
	(1.77)	(2.92)	(1.72)	(4.62)	(3.66)
Black × Unemployment	0.00394***	0.00404***	0.00394***	0.00464***	0.00364***
	(25.49)	(25.6)	(25.52)	(29.27)	(28.55)
Female	0.00550***	0.00560***	0.00552***	0.00579***	0.00420***
	(10.18)	(10.37)	(10.21)	(11.48)	(9.66)
Female × Unemployment	0.0000746	0.0000857	0.0000748	0.000102	0.000312***
	(0.81)	(0.93)	(0.81)	(1.19)	(4.14)
R^2	0.0415	0.0426	0.0415	0.0365	0.0435
Observations	17,939,045	17,939,045	17,939,045	17,939,045	18,477,713

Notes: The table reports the results for equation (1). The dependent variable is a dummy variable for (un)employment status. First column shows the results for *Female* and *Race* in equation (1). 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%.

Table 1 displays the estimation results for the coefficients on race and gender, respectively, i.e. γ_1 through γ_4 of equation (1). Column 1 contains the main specification just described, and

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

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

our interest is on the interactions of group status with the aggregate unemployment rate. The coefficient on "Black × 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 for instance in a severe recession aggregate unemployment increases from 5 to 10 percent. Using average values for the other control variables, we have that 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 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. Moreover, since it is well established that labor force participation has strong cyclical components, 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 C). 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

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, 2005)
b value of unemployment	0.9	Hagedorn and Manovskii (2008)
N_1 pop. share of group 1	90%	White/Black share in the labor market
c vacancy creation cost	0.329	Aggregate labor market statistics
ν employer's bargaining power	0.703	Aggregate labor market statistics
π degree of hiring discrimination	1.4	Resume studies

Notes: Parameter values used in the model. Aggregate statistics used to calibrate parameters c and ν are the long-run unemployment rate of 6.2%, unemployment volatility of 1.4 percentage points (both from the CPS), and an average job-finding rate of 45% (as in Shimer, 2005).

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, its standard deviation over time, as well as the 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 (2005), we use a separation rate of $s = 3.4\%$, and we follow Hagedorn and Manovskii (2008) in setting the value of unemployment to 0.9. Finally, in the CPS the ratio of blacks to whites among labor force participants is roughly 1 to 9.

We then pick the vacancy creation cost c and the employer's bargaining power ν in a way that minimizes the distance of the model outcomes to aggregate moments of the labor market in

Table 4: Model results

Outcome	Aggregate	Group 1	Group 2
Mean unemployment	6.74 %	6.65 %	7.50 %
Std dev unemployment	0.73 pp	0.71 pp	0.92 pp
Mean job-finding rate	47.1 %	47.7 %	41.9 %
Std dev job-finding rate	5.47 pp	5.46 pp	5.63 pp

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

the data. As these moments we use the long-run average level of unemployment and its standard deviation in the CPS data we use, which are 6.2% and 1.4%, respectively, and an average monthly job-finding rate of 45% as in Shimer (2005). The resulting parameter values are $c = 0.33$ and $\nu = 0.70$. This implies that employers receive around 70% of the match surplus. The creation of a vacancy costs an employer roughly 10 days worth' of output, and on aggregate hiring costs are just above 1% of average GDP. These parameter values are therefore well within the range frequently used in the literature. Table 4 shows that the model moments are reasonably close to the targets.

Finally, we pick the degree of hiring discrimination to be $\pi = 1.4$ for our baseline calibration, which by our count is the median estimate (in Jacquemet and Yannelis, 2012) of the resume studies surveyed in Baert (2017) 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 it has a first-order effect on the *distribution* of labor market outcomes between groups.[8]

Table 4 displays the model's results in terms of aggregate outcomes, and outcomes by groups. Notably, table 4 shows that there are both strong level effects and strong cyclical effects. 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

[8]The reason why there is a minimal effect on aggregate outcomes is that, with 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. In the case of full bargaining power by the employer there is no aggregate effect as wages are not affected (as in the example of section 2).

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 unemployment rises by 4% the unemployment rate for blacks increases close to 1.6 percentage points stronger than for whites. The model here implies that, in such a recession, the difference in unemployment rates increases by 1.17 percentage points, thus accounting for about 70% of our empirically measured gap.

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 difference in unemployment volatilities corresponds to our empirically measured value – this is the case at a value of $\pi = 1.55$. This is a value close to Bertrand and Mullainathan (2004)'s estimate of 1.5 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 volatility of unemployment rates (conditional on the model being correct). Increasing π in this way implies that in the model hiring discrimination now accounts for a steady-state difference in unemployment rates of 1.13 percentage points (compared to 0.85 points in the baseline), and for a stead-state difference in job-finding rates of 7.5 percentage points (compared to 5.8 points before).

Conversely, we can ask, how of the extra unemployment of blacks would be reduced if we could, say, cut the amount of hiring discrimination in half? We therefore set $\pi = 1.2$. In the model this means a reduction of black steady-state unemployment of 0.36 percentage points, and an increase in steady-state job-finding rates of 2.3 percentage points. Unemployment volatility is reduced and now exceeds the one of whites by 16%, instead of 31% in the baseline calibration[9]. This reduction in volatility means that in the case of our exemplary big recession with 4% higher aggregate unemployment, the black unemployment rate is an additional 0.54 percentage points lower than in a recession at current discrimination levels (the increase is now only 0.63 percentage

[9]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

points stronger than for whites', compared to 1.17 percentage points in the baseline).

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

Baert, S. (2017). Hiring Discrimination: An Overview of (Almost) All Correspondence Experiments Since 2005. *GLO Discussion Paper* (61).

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.

Black, S. E. and P. E. Strahan (2001). The division of spoils: Rent-sharing and discrimination in a regulated industry. *American Economic Review 91*(4), 814–831.

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

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).

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

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

Hagedorn, M. and I. Manovskii (2008). The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. *American Economic Review 98*(4), 1692–1702.

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 (2012). Indiscriminate discrimination: A correspondence test for ethnic homophily in the Chicago labor market. *Labour Economics 19*(6), 824–832.

Jarosch, G. and L. Pilossoph (2016). Statistical Discrimination and Duration Dependence in the Job Finding Rate.

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.

Neumark, D. (2016). Experimental Research on Labor Market Discrimination. (NBER working paper 22022), 1689–1699.

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

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

Appendix A Equilibrium conditions

To be added.

Appendix B Effects of parameters c and s

Appendix C Additional Robustness Checks

To be added.

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

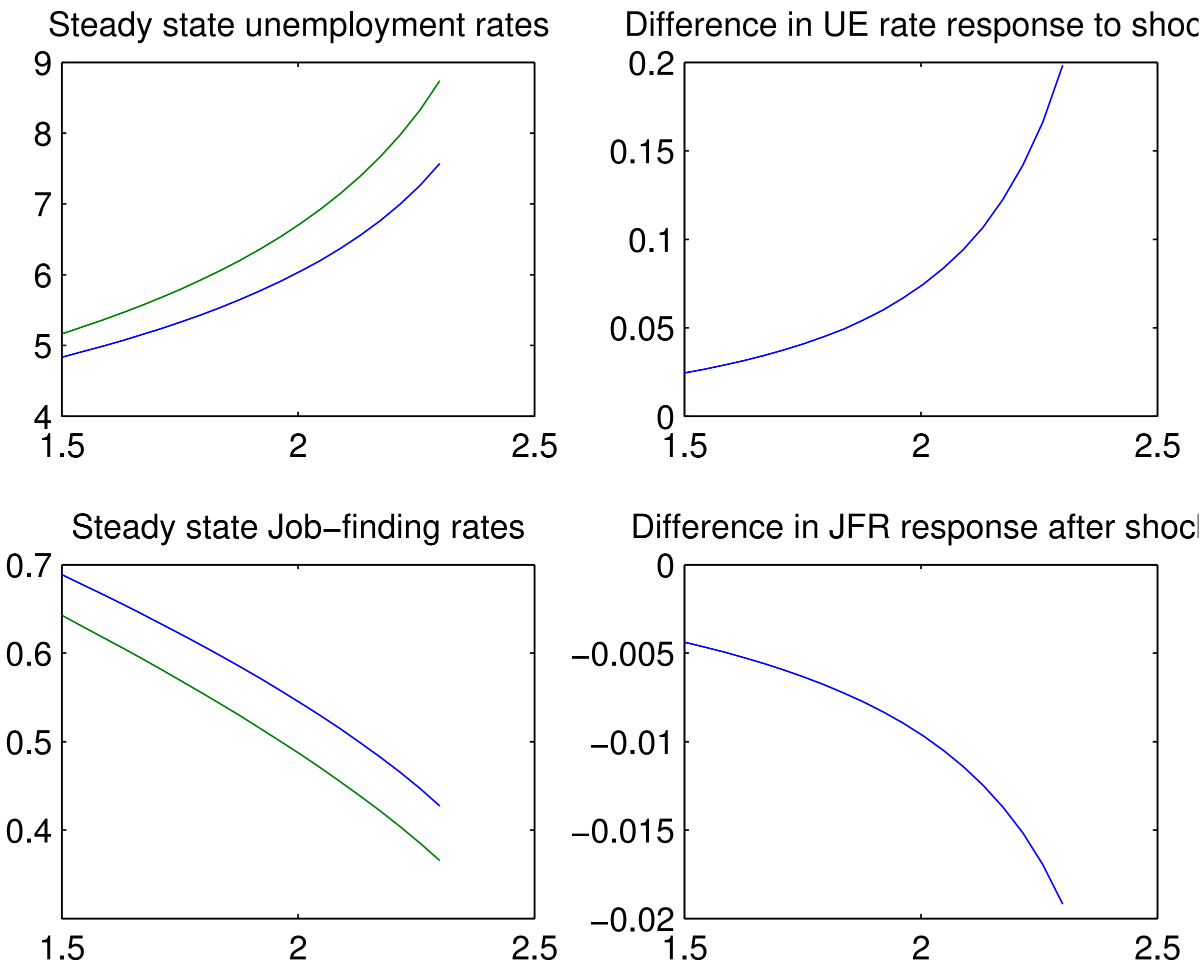

Notes: Model results for different values for c (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity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

Figure 4: Group differences as function of separation rate c

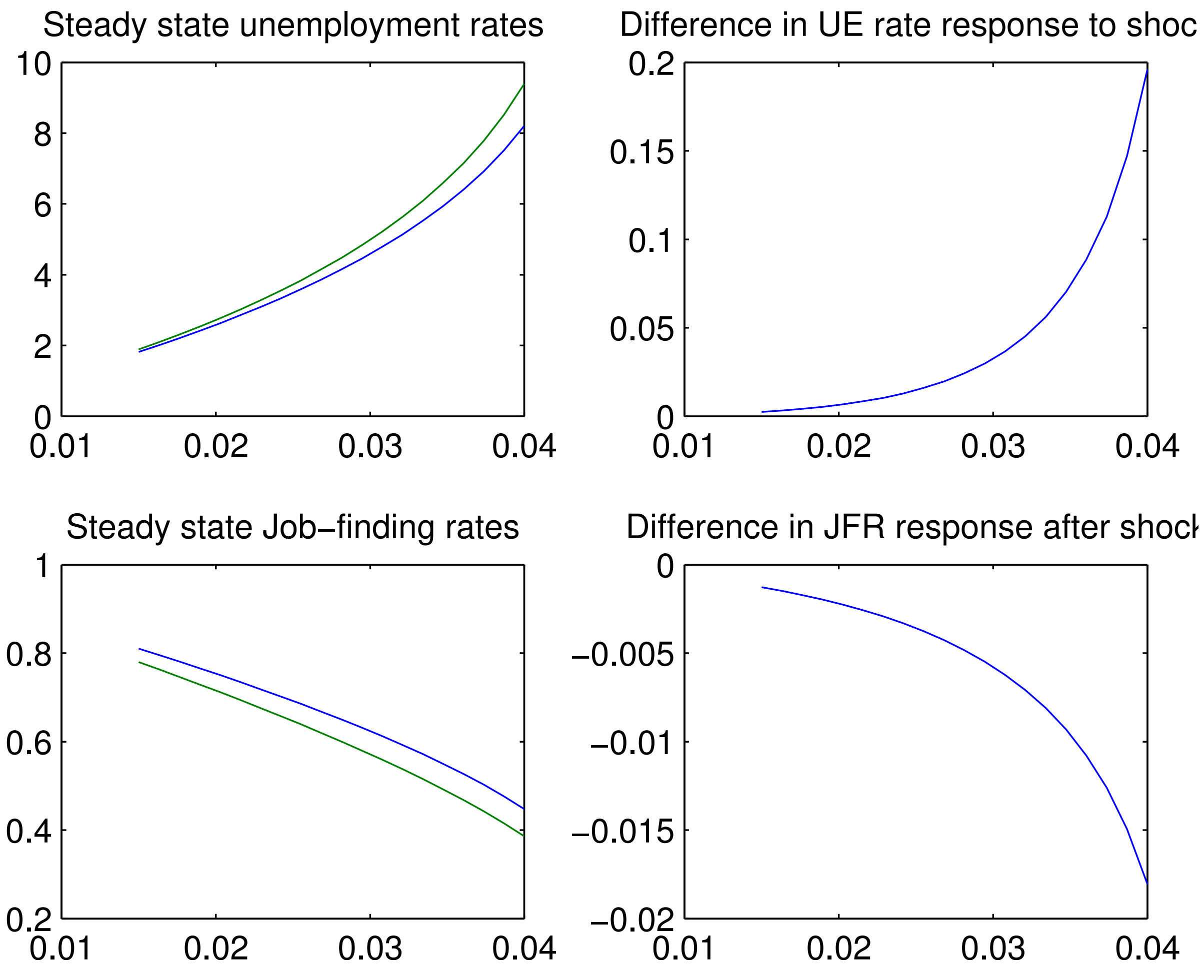

Notes: Model results for different values for c (x-axis). Left hand side panels: Steady-state unemployment rates (top left, in percent) and job-finding rates (bottom left). Right hand side panels: Difference (group 1 - group 2, in percentage points) in impulse responses after an increase in match productivity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