RACIAL DISCRIMINATION IN THE U.S. LABOR MARKET: EMPLOYMENT AND WAGE DIFFERENTIALS BY SKILL

Daniel Borowczyk-Martins
Jake Bradley
Linas Tarasonis

Discussion Paper 14 / 637

27 May 2014

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
United Kingdom

Racial Discrimination in the U.S. Labor Market: Employment and Wage Differentials by Skill

Daniel Borowczyk-Martins
University of Bristol

Jake Bradley
University of Bristol

Linas Tarasonis*
Aix-Marseille University (Aix-Marseille School of Economics), CNRS & EHESS

April 2014

Abstract

In the US labor market the average black worker is exposed to a lower employment rate and earns a lower wage compared to his white counterpart. Lang and Lehmann (2012) argue that these mean differences mask substantial heterogeneity along the distribution of workers' skill. In particular, they argue that black-white wage and employment gaps are smaller for high-skill workers. In this paper we show that a model of employer taste-based discrimination in a labor market characterized by search frictions and skill complementarities in production can replicate these regularities. We estimate the model with US data using methods of indirect inference. Our quantitative results portray the degree of employer prejudice in the US labor market as being strong and widespread, and provide evidence of an important skill gap between black and white workers. We use the model to undertake a structural decomposition and conclude that discrimination resulting from employer prejudice is quantitatively more important than skill differences to explain wage and employment gaps. In the final section of the paper we conduct a number of counterfactual experiments to assess the effectiveness of different policy approaches aimed at reducing racial differences in labor market outcomes.

Keywords: employment and wage differentials, discrimination, job search.

JEL codes: J31; J64; J71.

*The authors are indebted to Grégory Jolivet, Fabien Postel-Vinay and Hélène Turon for their encouragement and very insightful comments and suggestions. We also thank Pierre Cahuc, François Fontaine, Pieter Gautier, Thomas Jørgensen, Leo Kaas, Etienne Lalé, Jean-Marc Robin and various seminar audiences. Daniel Borowczyk-Martins acknowledges financial support from Fundação para a Ciência e a Tecnologia under grant SFRH/BD/38968/2007, co-funded by the European Social Fund. Part of the research for this paper was carried out while Linas Tarasonis was at CREST (INSEE) and at the University of Bristol.

1 Introduction

In their recent survey of the economic literature on racial discrimination Lang and Lehmann (2012) document persistent differences in employment and wages across black and white workers in the United States (US) labor market. They argue that negative black-white employment and wage differentials are the two main empirical regularities a model of discrimination should seek to replicate. Critically, their review of the evidence suggests these differentials vary considerably by skill. In particular, wage gaps '*are smaller or nonexistent for very high-skill workers*' and employment gaps are '*somewhat smaller among high-skill than among low-skill workers*' (p.12). The authors also assess the ability of existing discrimination models to replicate these facts. They conclude that '*no existing* [discrimination] *model can fully explain these regularities*' (idem).[1]

In this paper we develop a model of discrimination that successfully replicates these regularities. We show that a model of taste-based employer discrimination can deliver simultaneously mean black-white wage and employment gaps, as well as a decreasing profile of these gaps as the skill of workers increases. Recent evidence in the economics literature suggests labor market discrimination is still a plausible hypothesis to rationalize observed wage and employment gaps across races. Correspondence and audit studies find pervasive evidence of unequal treatment of black workers vis-a-vis seemingly equally skilled white workers (see Bertrand and Mullainathan (2004) and Charles and Guryan (2011)). The latest evidence produced by regression-based studies using the methodology proposed by Neal and Johnson (1996) points to the conclusion that, although differences in premarket factors are likely to play a major role in explaining observed mean black-white wage gaps, a substantial wage gap remains after controlling for premarket factors (see Carneiro et al. (2005) and Lang and Lehmann (2012)). Ritter and Taylor (2011) use Neal and Johnson's methodology to measure the importance of premarket factors for observed mean black-white employment gaps and also find that a substantial gap remains after controlling for premarket factors and a number of other variables.[2]

Two competing approaches dominate the economic literature on discrimination: prejudice or taste-based models, pioneered by Becker (1971), and models of statistical discrimination, starting with Phelps (1972) and Arrow (1973).[3] Despite the voluminious empirical literature on racial discrimination, there is no systematic evidence pointing to one approach as being more plausible than the other (see Charles and Guryan (2011) and Lang and Lehmann (2012)). We find both

[1] In this literature wage differentials are defined as one minus the ratio of mean black to white wages, whereas employment differentials refer to the percentage point difference between mean white and black employment rates.

[2] Neal's and Johnson's investigation of the role of premarket factors in explaining mean black-white wage gaps had a major impact in the literature. In that study the authors argue that controlling for differences in premarket factors (measured by the Armed Forces Qualification Test at an individual's young age) can fully explain the mean wage gap across blacks and whites in the US. As mentioned in the main text, recent evidence suggests a more nuanced view of the importance of premarket factors.

[3] The literature that followed Phelps (1972) focuses on the possibility that blacks' productivity is more difficult to observe than that of whites, while the literature that builds on Arrow (1973) stresses the effects of differences in employers' beliefs about blacks' and whites' productivity.

approaches compelling on theoretical grounds and believe both are important to rationalize the data. However, in this paper we set aside statistical discrimination and focus only on the consequences of taste-based discrimination for differences in labor market outcomes of blacks and whites. In particular, we build on a modeling approach that combines employer taste-based discrimination and random search frictions to describe differences in labor market outcomes of workers who differ in terms of a nonproductive attribute (e.g. race or gender). As documented by Lang and Lehmann (2012), among existing models this approach is the only one that can predict both mean employment and wage gaps. The assumption of search frictions is a natural modeling choice in this context, since there are sizeable differences in mean unemployment durations of black and white individuals.

Previous models in the discrimination literature have shown how the combination of taste-based discrimination and random search frictions generates mean employment and wage differentials across races.[4] The central intuition is the following. Consider an economy populated by two types of workers (who differ by race) and two types of employers, where one type (prejudiced) incurs a utility cost from hiring a black worker. Since prejudice reduces the match value between prejudiced employers and black workers, the matching opportunities of black workers are smaller compared to those of white workers. Under random search, black workers cannot direct their search away from prejudiced employers (their probability of meeting a prejudiced employer is the same as that of whites), so in this setup black workers have lower employment prospects compared to white workers. This delivers mean employment differentials across races. Lower employment prospects in turn imply black workers have lower reservations values. Since all employers know it takes longer for black workers to find a job, they will take advantage of that and offer black workers lower wages. This delivers mean racial wage differentials.

To generate wage and employment gaps that are smaller for high-skill workers compared to low-skill workers we extend this modeling approach in two directions. We start by assuming that workers and firms differ respectively in their levels of skill and technology and that the production value of the match is a complementary function of firms' and workers' skill levels. Under production complementarities and capacity constraints, in equilibrium, high(low)-skill workers will be matched more frequently with high(low)-technology firms. In other words, there will be positive assortative matching on skill. Because production is also an increasing function of workers' and firms' skill levels, matches involving high-skill workers will involve higher levels of production. If we further assume the utility cost for prejudiced employers of employing a black worker is constant, then the cost of prejudice represents a smaller share of the value of matches as the skill of workers increases. This will translate into lower employment and wage differentials across races as the level of workers' skill increases. To see this more clearly, consider a black worker with a low level of skill. Due to assortative matching in skill, only low-technology firms

[4] This result was first shown in Black (1995).

will be willing to match with him. On the other hand, the utility cost of prejudice represents a high share of the match value between low-skill workers and low-technology firms, which implies that few of these matches are viable. Therefore, in equilibrium, low-skill black workers' employment prospects are considerably worse compared to their white counterparts, and since lower reservation values feed into lower wages, their wages are also considerably smaller compared to whites. Now consider a high-skill black worker. Since he matches with high-technology firms, prejudice will play a negligible role in those firms' hiring decisions and, by extension, to their wage policies across races, which implies smaller wage and employment differentials across races for high-skill workers.

The model developed in this paper shares several features with other models of taste-based discrimination in a random search environment, like Black (1995), Bowlus and Eckstein (2002), Rosén (2003) and Flabbi (2010).[5] In all these models prejudiced employers incur a psychic cost of employing a worker who belongs to the minority group. Black (1995) assumes prejudiced employers never match with black workers. Bowlus and Eckstein (2002) develop a wage posting model and allow both types of worker to draw their productivity from separate (degenerate) distributions. Rosén (2003) and Flabbi (2010) model the productivity of the match (and not of workers and jobs) as being heterogeneous.

Our model setup differs from these papers by assuming two-sided skill heterogeneity, production complementarities and endogenous vacancy posting. These assumptions have strong empirical support. The analysis of matched employer-employee data sets over the past 15 years has consistently documented the importance of both worker and firm unobserved heterogeneity to explain observed differences in wages across workers (see Abowd et al. (1999), Lopes De Melo (2013) and Torres et al. (2013)). Production complementarities are increasingly seen as a plausible description of the production technology characteristic of modern labor markets. Shimer and Smith (2000) established the equivalence between a specific form of production complementarities and positive assortative matching in a random search environment, generalizing the famous result in a competitive setup put forth by Becker (1973). Several papers in the applied search literature have tried to measure the sign and strength of assortative matching in skill nonparametrically (see Abowd et al. (1999), Eeckhout and Kircher (2011), Hagedorn et al. (2012), Torres et al. (2013) and Lopes De Melo (2013)), or estimate the degree of skill complementarities in production using a structural approach (see Bagger and Lentz (2012) and Lise et al. (2013)). Although this question is not fully settled, there is a growing consensus that positive sorting on skill is an important characteristic of modern labor markets. Finally, the assumption of endogenous vacancy posting is motivated by the vast evidence supporting the existence of an aggregate matching function in the US labor market and firms' response to

[5] We only refer to papers based on a similar modeling approach. See Lang and Lehmann (2012) for a comprehensive review of all discrimination models.

changes in vacancy filling rates (see Petrongolo and Pissarides (2001), Borowczyk-Martins et al. (2013) and Davis et al. (2013)). The importance of this assumption surfaces when we use the model to conduct counterfactual analysis, as it allow us to take into account general equilibrium effects of different policy approaches.

We estimate the model using various sources of publicly available data for the US manufacturing sector and indirect inference estimation methods.[6] A critical feature the model must satisfy to make its empirical implementation plausible is to allow for the possibility that black and white workers have different skill distributions. Indeed, there is substantive evidence of persistent black-white gaps in educational attainment and cognitive skill (see Neal (2006)), which suggests differences in skill across races are likely to play an important role in shaping mean employment and wage differentials. Other models of taste-based employer discrimination have been estimated using structural methods (see Bowlus and Eckstein (2002) and Flabbi (2010)). However, our paper is the first to take to the data a search-discrimination model based on Shimer and Smith's (2000) partnership model. Structural estimation of this vintage of models is very recent in the applied search literature (see Jacquemet and Robin (2012) and Lise et al. (2013)). An important distinction of our estimation strategy with respect to other structural papers in the literature is that it allows prejudice to generate racial differences not only in wages, but also in unemployment rates.

We estimate that about half the employers (49%) are prejudiced against black workers and that the utility cost of employing a black worker is about 10.5% of the average productivity of a match involving white workers. These results portray the degree of employer prejudice in the US labor market as being strong and widespread. We also find differences in the skill distributions of black and white individuals. The mean skill of black workers is estimated to be 4.95% lower than whites'. We use the estimated model to address a number of questions that have been central to the literature. We characterize the equilibrium matching patterns between workers of different race and skill and jobs with different technology and operated by employers with different racial attitudes. The second contribution we offer is a precise, model-based, quantitative account of the sources underlying the empirical regularities we seek to replicate. We decompose the observed black-white differences in wages and employment into the contribution of employer prejudice and differences in skill across races. This is an important undertaking as the literature identifies discrimination and skill differences as the two main competing explanations for the existence of differences in labor market outcomes across races. We find that, to describe the differences in labor market outcomes of black and white males in the US, both skill differences across races and employer prejudice are required, but discrimination generated by prejudiced employers is quantitatively more important. According to our structural decomposition, employer prejudice

[6] Two-sided skill heterogeneity is a distinctive feature of the model which calls for an estimation strategy that makes use of matched employer-employee data. Unfortunately, for the US labor market no comprehensive database is publicly available.

explains almost the full observed mean unemployment gap across races, while skill differences are quantitatively more important to explain observed differences in the top quantiles of the wage distributions.

In the last section of the paper we conduct a number of counterfactual experiments to gauge the effectiveness of alternative policy approaches to improve labor market outcomes of black vis-a-vis white workers. We consider three policy approaches: premarket affirmative action, labor market affirmative action and law enforcement, and evaluate their impact both on social welfare (efficiency) and the individual welfare of workers and firms of different types and skill levels (redistribution). Our results seem to suggest policy can improve labor market outcomes of black workers while simultaneously increasing social welfare. We discuss the limitations of our analysis and how they may affect our conclusions.

The paper is structured in the following way. In Section 2 we present the theoretical model, derive its equilibrium and theoretical properties. Section 3 describes the data. The estimation procedure is presented in detail in Section 4. Section 5 describes the fit of the model, while Section 6 reports the structural parameter estimates. In the remaining sections we explore several applications using the estimated model. Section 7 analyzes the equilibrium sorting patterns of our model economy, Section 8 quantifies the relative importance of prejudice and skill differences in explaining differences in black-white employment and wage gaps, while Section 9 studies the effects of alternative policy approaches. Section 10 concludes.

2 The Model

The model we develop in this section builds on Shimer and Smith's (2000) partnership model, extending it to a labor market where some employers are prejudiced vis-a-vis a specific type of workers and in which there is free entry of jobs. Because the model applies to any market where some employers are prejudiced against a certain type of worker, we will adopt a more general terminology in this section and return to the racial discrimination application we have in mind in the estimation section. Section 2.6 contains the main results of interest. It describes the properties of the equilibrium in the extended model — what we call a *dual sorting equilibrium*. The sections that precede Section 2.6 set out the model in detail and discuss the relevance of the model's assumptions to study discrimination in the labor market.

2.1 The Environment

We consider a labor market with L workers and G jobs. The number of jobs G will be determined in equilibrium and L is given. There exist two types of employers and two types of workers. A share m of workers are of type-1 and a share $(1-m)$ of them are of type-2, with worker types being denoted by index $i = 1, 2$. Similarly, a share π of jobs are operated by prejudiced employers (P) and a share $(1-\pi)$ of those are not (N), with index $j = P, N$ denoting respectively jobs operated by prejudiced and nonprejudiced employers. In this model one firm is one job. Thus, throughout the text we use the terms *jobs*, *employers* and *firms* interchangeably. Workers further differ in their skill h, which we assume is uniformly distributed over the unit interval, $h \sim U(0,1)$, and firms in their level of technology (efficiency of labor inputs) x, which we also assume is uniformly distributed in the unit interval, $x \sim U(0,1)$.[7] Let $\ell_i(h)$ and $g^j(x)$ denote respectively the population measures of type-i workers of skill h and type-j firms of technology x.[8] The endogenous measures of type-i unemployed workers of skill h and type-j vacant firms of technology x are respectively denoted $u_i(h)$ and $v^j(x)$, with total measures of type-i unemployed workers and type-j vacant jobs given respectively by $u_i = \int u_i(h)\,\mathrm{d}h$ and $v^j = \int v^j(x)\,\mathrm{d}x$.

Time is continuous and both workers and employers are risk neutral, with discount rate ρ. Employers and workers maximize the present discounted value of future utility streams, measured in monetary terms. As in Becker (1971), prejudiced employers incur a psychic cost d of employing a type-2 worker. When a worker and an employer meet, their flow output depends on their levels of skill, denoted $f(h,x)$ and satisfying certain regularity conditions (see Appendix A). We take complementarities in skill as a descriptive feature of modern labor markets and so assume a supermodular production function.[9] This means the own marginal product of any worker and job is increasing in his partner's skill. Formal details are provided in Appendix A.[10]

We assume that only unemployed workers and vacant firms search for a partner, ruling out on-the-job search.[11] All unemployed workers search for jobs with equal search intensity. Job offers and unemployed job applicants arrive respectively to unemployed workers and vacant firms following a Poisson process. At each point in time the job and unemployed arrival rates are a function of the number of searchers on each side of the market via the aggregate meeting function $M(u_1+u_2, v^P+v^N)$. In the meeting process, type-1 and type-2 workers, and type-N and type-P firms, are perfect substitutes. Meeting is random and vacant jobs and unemployed workers of different types and skills are effectively exposed to the same arrival rates, respectively $\lambda^W = \frac{M(u_1+u_2,v^P+v^N)}{u_1+u_2}$ for workers, and $\lambda^F = \frac{M(u_1+u_2,v^P+v^N)}{v^P+v^N}$ for jobs.

Once a firm and a worker meet they decide whether or not to form a match. We define a match indicator function $\alpha_i^j(h,x)$, which is equal to 1 if a type-i worker of skill h and a type-j

[7] These assumptions about the supports and densities of the skill distributions of workers and firms are normalizations. In the empirical application we allow all skill distributions to be defined over distinct supports and their densities to be non-uniform. Since we assume the production function is increasing in worker's and firms' skill, one can think of h and x as the skill ranks of the underlying skill distributions.

[8] So that $mL = \int \ell_1(h)\,\mathrm{d}h$ and $\pi G = \int g^P(x)\,\mathrm{d}x$.

[9] There is a growing consensus in the applied search literature on the descriptive relevance of positive assortative matching (see Lise et al. (2013), Lopes De Melo (2013) and Bagger and Lentz (2012)).

[10] In our empirical application we assume a specific degree of skill complementarities. However, the model, as well as the equilibrium characterization described in Section 2.6, hold with any production function that satisfies the regularity conditions and supermodularity.

[11] Empirically, we observe mean differences in job-to-job transition rates across races. Introducing on-the-job search to account for this fact would add significant complexity to the model without improving substantially the quantitative performance of the model in terms of replicating cross-sectional differences in wages and unemployment rates across races.

firm of technology x decide to match upon meeting. Matches are randomly destroyed by a Poisson process with arrival rate δ, in which case both the worker and the firm join the pool of searchers. This job destruction rate is assumed constant, irrespective of firm and worker type. We discuss the quantitative implications of this assumption in Section 4.2.1.

2.2 Value Functions

A worker of type-i and skill level h can be in one of two different states: employed or unemployed. The flow value of employment for a type-i worker of skill h employed in a type-j firm of technology x is given by equation (1), where $w_i^j(h,x)$ is the wage she earns in this job and $U_i(h)$ is the value of being unemployed. The wage $w_i^j(h,x)$ is endogenously determined as the solution to a Nash bargaining game.

$$\rho W_i^j(h,x) = w_i^j(h,x) + \delta \left[U_i(h) - W_i^j(h,x)\right]. \tag{1}$$

While unemployed a worker receives a flow utility b, independent of his race and skill. Therefore, the value of being unemployed for a worker of type-i and skill h, $U_i(h)$, independent of the worker's employment history, is given by the following equation:

$$\rho U_i(h) = b + \lambda^W \sum_{j=P,N} \int \alpha_i^j(h,x) \left[W_i^j(h,x) - U_i(h)\right] \frac{v^j(x)}{v^P + v^N} \mathrm{d}x. \tag{2}$$

Note that $\frac{v^j(x)}{v^P+v^N}$ is the probability of sampling a vacant type-j job of technology x from the pool of unmatched firms.

Firms can be in two states: vacant or filled. $J_i^j(h,x)$ is the value of a job for a type-j firm of technology x, filled with a type-i worker of skill h:

$$\rho J_i^j(h,x) = f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - w_i^j(h,x) + \delta \left[V^j(x) - J_i^j(h,x)\right], \tag{3}$$

where the indicator function, $\mathbf{1}_{[(i,j)=(2,P)]}$ takes value 1 if the match involves a prejudiced employer and a type-2 worker.

In our model d is a psychic cost valued in monetary terms that a prejudiced employer incurs upon matching with a type-2 worker. This psychic cost does not affect the production value of the match but it does affect the utility value of the match for a prejudiced employer. Note that in our specification d enters additively in the firm's value equation.[12] This implies the degree of prejudice is independent of workers' and firms' skill levels. From a descriptive point of view, there is no reason to think this specification is less plausible than other alternatives (e.g. assuming the level of prejudice is decreasing in firm's level of technology). To replicate the key stylized fact that provides the main motivation for our model we could allow for some heterogeneity in the disutility of prejudice, as long as the level of prejudice is not exactly proportional to the

[12] This assumption is standard in search models of taste-based discrimination with bargaining and match-specific heterogeneity (see Rosén (2003) and Flabbi (2010)), as well as in search models with wage posting (see Bowlus and Eckstein (2002)).

match production value. However, from an empirical point of view, a constant utility cost of prejudice is important to secure identification.

We assume that posting a vacancy has a nonnegative flow cost $\kappa \geq 0$. The value of posting a vacancy for a type-j firm of technology x, $V^j(x)$, depends on the probability of the vacancy being filled by each of the two types of worker and it is given by the following equation:

$$\rho V^j(x) = -\kappa + \lambda^F \sum_{i=1,2} \int \alpha_i^j(h,x) \left[J_i^j(h,x) - V^j(x) \right] \frac{u_i(h)}{u_1+u_2} \mathrm{d}h, \tag{4}$$

where $\frac{u_i(h)}{u_1+u_2}$ is the probability of sampling an unemployed type-i worker of skill h from the pool of unmatched workers.

2.3 Entry

We assume that jobs remain active in the market if the present discounted value of keeping a job unfilled is nonnegative, i.e. if $V^j(x) \geq 0$. To determine the total mass of active jobs in equilibrium, G, we assume free entry. We show in Proposition 2 in Appendix A that $V^j(x)$ is strictly increasing in x. We can therefore make a normalization and assume the least efficient job operated by a nonprejudiced employer makes zero profit. Free entry of jobs implies the following conditions respectively for jobs operated by nonprejudiced and prejudiced employers:

$$V^N(0) = 0, \tag{5}$$

$$V^P\left(x^{P\star}\right) = 0, \tag{6}$$

where $x^{P\star}$ is the technology level of the least efficient job operated by a prejudiced employer. We show in Corollary 4 that, if there is employer prejudice in this economy, i.e. if $\pi \in (0,1)$ and $d > 0$, then $x^{P\star} > 0$. That is, if the underlying technology distribution is common to both employer types, the least efficient job operated by a prejudiced employer in the market has a higher technology level compared to the least efficient job operated by a nonprejudiced employer.[13]

Prior to entry the characteristics of the employer who will operate the job and the technology level of the job are unknown. With probability π the employer is prejudiced. The technology level of new jobs is drawn from the probability distribution, $g(x)$. As the mass of jobs increases, the expected value of keeping the job unfilled decreases for all active jobs. This process stops when the expected value of keeping a job unfilled is zero for the least efficient jobs operated by prejudiced and nonprejudiced employers. One possible interpretation of this model of entry in the context of labor market discrimination is the following. At any point in time an economy generates a certain number of jobs. These jobs differ in their efficiency of labor usage and on

[13] The normalization we make here implicitly assumes $d \geq 0$. We only focus on equilibria satisfying this condition.

the racial attitude of the individual who is responsible for hiring a worker to operate it and who captures the surplus generated by the job. Because prejudice is costly, in equilibrium, the least efficient jobs operated by prejudiced employers have to be more efficient compared to the least efficient jobs operated by nonprejudiced employers. A prejudiced employer who draws a job with technology level lower than $x^{P\star}$ immediately leaves the market. Since the threshold technology level required for entry differs depending on a firm's racial attitude (see Corollary 4), the distribution of productivity amongst active prejudiced firms will differ from that of active nonprejudiced firms[14].

A question that has received a great deal of attention in this literature concerns the survival of prejudiced firms in the long run equilibrium of the economy (see e.g. Arrow (1973) and Rosén (2003)). Our model of entry resembles other descriptions in the literature. Black (1995) makes a similar argument to ours, but in his model prejudiced firms require a higher draw of entrepreneurial ability to enter the market, where the latter is independent of output. In the model developed by Rosén (2003) survival of prejudiced firms is achieved by a separation between owners and managers, where the owners are the residual claimants on output and managers bear the cost of prejudice. In the context of this model, prejudice firms can exist, but must be more productive on average in order to do so.

2.4 Match Surplus

From the four value functions written above, we can determine the total surplus generated by any match. The surplus of a match between a type-i worker of skill h employed in a type-j firm of technology x is $S_i^j(h,x) = W_i^j(h,x) - U_i(h) + J_i^j(h,x) - V^j(x)$ and it is split between the worker and firm according to Nash bargainng. The worker takes a share β and the firm a share $(1-\beta)$, implying the following equalities:

$$S_i^j(h,x) = \frac{J_i^j(h,x) - V^j(x)}{1-\beta} = \frac{W_i^j(h,x) - U_i(h)}{\beta}. \tag{7}$$

The wage equation that solves this bargaining problem is given by the following expression:

$$w_i^j(h,x) = \beta\left[f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \rho V^j(x)\right] + (1-\beta)\rho U_i(h). \tag{8}$$

Using equations (1) and (3), the expression for the total surplus can be rearranged and expressed in the following way:

$$S_i^j(h,x) = \frac{f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \rho U_i(h) - \rho V^j(x)}{\rho+\delta}. \tag{9}$$

Whenever the surplus is positive a match is formed. Formally, we denote this by an indicator function:

$$\alpha_i^j(h,x) = \mathbf{1}\left[f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \rho U_i(h) - \rho V^j(x) > 0\right]. \tag{10}$$

[14] Formally, $\frac{g^P(x)}{\pi G} \neq \frac{g^N(x)}{(1-\pi)G}$.

A first remark about the bargaining process is that we assume firms and workers of different types have the same rent-sharing parameter. Allowing for different rent-sharing parameters across worker types would introduce another degree of heterogeneity across types. We choose not to follow this route. Instead, we take the view of the strategic bargaining literature that rent-sharing parameters measure the relative impatience of bargaining participants and see no reason for it to differ across worker types and/or skill levels.[15] Moreover, in the context of our model, absent taste-based discrimination (if d were zero) a lower β for type-2 workers by itself would not generate employment differences.

A second remark concerns the fact that d is transferable among match partners. This means the psychic cost of prejudiced firms is observable by both parties in the match and shared among them according to their rent-sharing parameters. We acknowledge that alternative specifications (for instance, one in which d is nontransferable) may be equally plausible. In Section 9 we simulate a model in which d is not transferable and analyze its effects on labor market outcomes and individual welfare. It should be noted that, although these two specifications tell a somewhat different story about how prejudice translates into discrimination, they have similar properties. In particular, both generate hiring and wage discrimination, as well as higher employment and wage gaps for low-skill workers compared to high-skill workers.

Having established the structure of agents' payoffs, we can now define each agent's strategy. For a type-i worker of skill h her strategy is given by two sets, $\mathcal{M}_i^P(h)$ and $\mathcal{M}_i^N(h)$. Similarly, a firm's strategy is defined by two sets, $\mathcal{M}_1^j(x)$ and $\mathcal{M}_2^j(x)$. An agent's matching sets contains all the acceptable partners with whom she is willing to match and who are willing to match with her. The symmetry of matching sets is due to the surplus-sharing rule being jointly privately efficient (i.e. the decision to match is mutually agreeable). Using the indicator function $\alpha_i^j(h,x)$ we can express each worker's matching set as:

$$\mathcal{M}_i^j(h) = \{x \mid \alpha_i^j(h,x) = 1\}, \tag{11}$$

and each firm's matching set as:

$$\mathcal{M}_i^j(x) = \{h \mid \alpha_i^j(h,x) = 1\}. \tag{12}$$

2.5 Steady-state Equilibrium

A steady-state equilibrium in this model is characterized by four conditions: (i) workers and firms maximize their expected payoff, taking the strategies of all other agents as given; (ii) agents decide to match if it increases their payoff; (iii) all measures of unmatched workers of type-i and skill h and firms of type-j and technology x, resp. $u_i(h)$ and $v^j(x)$, are in steady-state and (iv) the least efficient active firms make zero profit. Conditions (i) and (ii) are given respectively by

[15] Some papers interpret differences in rent-sharing parameters as the result of discrimination (see Eckstein and Wolpin (1999) and Bartolucci (Forthcoming)).

firms' and workers' value functions and their matching sets. Condition (iv) is given by the entry conditions. We now state the assumptions necessary to ensure condition (iii).

To fix the measures of unmatched agents, flow creation and flow destruction of matches for every type of agent must exactly balance. This is given by the following set of equations:

$$\lambda^W \alpha_i^j (h, x)\, u_i (h) \frac{v^j (x)}{v^P + v^N} = \delta \gamma_i^j (h, x)\,, \tag{13}$$

where $\gamma_i^j (h, x)$ is a joint measure of matched type-i workers of skill h and type-j firms of technology x. These equations ensure that, for every possible match between a worker and firm of different types and skill levels, the number of matches being created at every point in time (the left-hand side of equation (13)) is exactly the same as the number of matches being destroyed (the right-hand side of equation (13)). Then, by definition, the steady-state stock of type-i employed workers of skill h is given by the following equation:

$$\ell_i (h) - u_i (h) = \int \gamma_i^N (h, x)\, \mathrm{d}x + \int \gamma_i^P (h, x)\, \mathrm{d}x. \tag{14}$$

That is, the total population of type-i workers of skill h must equal the sum of its unemployed and employed populations. Similarly, we can define the population of active type-j firms of technology x by:

$$g^j (x) - v^j (x) = \int \gamma_1^j (h, x)\, \mathrm{d}h + \int \gamma_2^j (h, x)\, \mathrm{d}h,\ \forall x\, :\, V^j(x) \geq 0, \tag{15}$$

so that, for each type-j firm of technology x, the total number of firms $g^j (x)$ is equal to the total number of matched and vacant firms. Inactive firms post zero vacancies $v^j (x) = 0$, $\forall x$: $V^j(x) < 0$.

To obtain the equilibrium conditions of the model we first use the bargaining solution (equation (7)) and firms' and workers' value functions to write each agent's equilibrium reservation value (equations (32) and (33)), where matching decisions satisfy equation (10). We then use the flow-balance equations (equation (13)) and the population accounting equations (equations (14) and 15) to express the equilibrium measures of unmatched agents (equations (34) and (35)). Finally, the number of firms and the truncation point of prejudiced firms' technology distribution is given by the two entry conditions (equations (5) and (6)). The formal definition of equilibrium is given in Appendix B.

2.6 Dual Sorting Equilibrium

We now explore some implications of equilibrium for workers and jobs of different types and skill levels. Proofs of stated results can be found in Appendix A. The results we present in this section are characteristic of a class of equilibria where, for all possible combinations of workers and firms of different types, some but not all matches are feasible. In other words, equilibria in which, for any combination between a type-i worker and a type-j firm, some matches between workers of skill h and firms of technology x are formed, but not for all possible combinations of skill and technology levels. In practice this implies that, on the one hand, the flow value of unemployment b, the psychic cost borne by prejudiced employers d and firms' vacancy cost κ are sufficiently small with respect to the value of production $f(h, x)$, so that, for all combinations of workers and firms of different types, there exist combinations of skill levels (h, x) that satisfy the match feasibility condition. On the other hand, it implies that, for certain combinations (h, x), the value of production $f(h, x)$ is small enough with respect to b, κ and d to render the match between them not feasible, where this holds for any combination (type-i, type-j).[16] In the various simulations of the model carried out in the execution of the paper we always found equilibria satisfying this description.

An equilibrium in our model economy is characterized by two forms of sorting across workers and jobs of different types and skill levels. To obtain positive assortative matching in skill we assume the production function is supermodular. We do not prove formally that, in our environment, assuming a supermodular production function implies positive assortative matching. However, for all the simulations of the model we performed using different ranges of parameter values, we always observe positive assortative matching in skill. Regarding the second form of sorting – the patterns of negative assortative matching between black/white workers and prejudiced/nonprejudiced employers – we are able to establish some analytical results.

Before proceeding with the statements of the model's implications we establish some necessary definitions. In the economics literature, discrimination is said to exist when equally productive workers are treated differently based on nonproductivity related factors, such as race or gender (see Cain (1986)).[17] In our model, the psychic cost d reduces the utility value of the match for prejudiced employers, but, importantly, it does not affect the production value of the match, so the way we model discrimination is broadly consistent with the traditional definition studied in the literature. The first instance of economic discrimination we are interested in characterizing pertains to workers' wages.

Definition 1 (Wage Discrimination): *A type-i worker of skill h experiences wage discrimination if she is paid a lower wage than an equally able type-k ≠ i worker when both are matched with type-j firms with the same technology x, that is*

$$\textit{for some } (h, x), \; w_i^j(h, x) < w_{k \neq i}^j(h, x).$$

The second instance of economic discrimination we characterize relates to agents decision of whom to match with. In an economy with no match surplus losses due to prejudice, type-1 and

[16] These two conditions can be stated formally in the following way: $\forall (i, j) \in [1, 2] \times [N, P]$, $\exists (h', x')$ and $(h'', x'') \in [0, 1]^2$ such that $\alpha_i^j(h', x') = 1$ and $\alpha_i^j(h'', x'') = 0$.

[17] As Cain (1986) emphasizes '*although physical productivity excludes the psychic component* [it] *is intended to be broad and to include such characteristics of the workers as their regularity in attendance at work, dependability, cooperation, expected future productivity with the firm, and so on.*'

type-2 workers with the same skill match with firms within the same range of technology. In an economy in which there is employer prejudice, in general, this will no longer be the case and the matching sets of two equally skilled workers of different types will differ. One reason why these matching sets differ is due to hiring discrimination. Formally, we have that:

Definition 2 (Hiring Discrimination): *A type-i worker of skill h experiences hiring discrimination if, upon meeting a firm of technology x of type-j, he is not hired, but an equally skilled type-$k \neq i$ worker is; that is,*

$$\text{for some } (h,x),\ \ \alpha_i^j(h,x) = 0 \ \text{ and } \alpha_{k\neq i}^j(h,x) = 1.$$

To fix ideas, note that hiring discrimination describes discriminatory behavior by employers that is materialized in the decision to hire a worker — a decision that is different from that of how much to pay him (wage discrimination), but that stems from the same cause, viz. prejudice. The first implication of our model is that, for a positive value of d, type-P firms (those who are prejudiced) and type-2 workers (those who are the object of prejudice) face worse prospects in the labor market compared to type-N firms and type-1 workers, respectively. This result is stated in the following proposition.

Proposition 1 (Outside Option Effects): *For any equilibrium such that $\pi \in (0,1)$ and $d > 0$:*

(i) for a worker of skill h, the value of unemployment of a type-1 worker is higher than that of a type-2 worker, that is, $U_1(h) > U_2(h)$, $\forall\, h$; and

(ii) for a firm of technology x, the value of a vacancy to a type-N firm is higher than to a type-P firm, that is, $V^N(x) > V^P(x)$, $\forall\, x$.

A corollary of Proposition 1 is that, if there are any prejudiced employers in this model economy, there will be wage discrimination against all type-2 workers.

Corollary 1 (Type-2 Wage Discrimination): *For any equilibrium such that $\pi \in (0,1)$ and $d > 0$, all type-2 workers experience wage discrimination in both types of firms and of any technological level.*

So far we have established that, for prejudiced employers, the decision to match with a type-2 worker differs from that of matching with an equally skilled type-1 worker in two ways. First, the psychic cost d directly reduces the utility value of the match with a type-2 worker. Second, type-2 workers have a lower outside option, which increases the value of the match. Since the decision to match is only governed by the match surplus condition (see equation (10)), matching between type-2 workers and type-P firms depends on the relative magnitude of these two effects.

We can show that, in our environment, the presence of employer prejudice implies that some type-2 workers will not be hired by certain prejudiced firms (the first effect dominates the second and it is high enough to reduce the match surplus to zero) and that those same firms will hire an equally able type-1 worker. As stated in Definition 2 this combination of circumstances entails hiring discrimination. The result is stated below.

Corollary 2 (Type-2 Hiring Discrimination in Prejudiced Firms): *For any equilibrium such that* $\pi \in (0,1)$ *and* $d > 0$*, some type-2 workers experience hiring discrimination by some prejudiced firms.*

For nonprejudiced firms, the difference between matching with equally skilled type-1 and type-2 workers is only affected by the outside option effect. When the match surplus condition between a nonprejudiced firm and a type-1 worker is not satisfied, the lower outside option of an equally skilled type-2 worker may render that match feasible. We will refer to this combination of circumstances as *reverse hiring discrimination*. This result is stated below.

Corollary 3 (Type-1 Hiring Discrimination in Nonprejudiced Firms): *For any equilibrium such that* $\pi \in (0,1)$ *and* $d > 0$*, some type-1 workers experience hiring discrimination by some nonprejudiced firms.*

We now turn our attention to the distributions of technology among prejudiced and nonprejudiced firms in equilibrium. If there are any type-2 workers in the economy, prejudiced firms have to be more efficient than nonprejudiced firms, since they have to make up for the cost of prejudice d — which affects negatively their matching opportunities and the match value of feasible matches with type-2 workers. This effect is stated in the following corollary.

Corollary 4 (Threshold Technology Differences): *For any equilibrium such that* $\pi \in (0,1)$ *and* $d > 0$*,* $\exists\, x^{P\star} > 0$ *such that* $V^P\left(x^{P\star}\right) = 0$*.*

Corollary 4 implies that, conditional on the technology distributions across firm types being the same, in equilibrium, prejudiced firms will be on average unambiguously more efficient than nonprejudiced firms.

3 Data

To estimate the model we use three sources of data: worker-level data, firm-level data and market-level data. In the following paragraphs we describe these data sources and the procedures we implemented to obtain the sample used for estimation. The moments of the estimation sample are presented and discussed in Section 5.

The worker-level data comes from the Current Population Survey (CPS). We merge the Monthly Outgoing Rotation Groups (MORG) with the Basic Monthly (BM) extracts, thus gathering information on individual wages and transition rates across employment and unemployment. Our sample runs from May 2004 to December 2005. We limit the sample to include only individuals who declare themselves to be either black or white. We only keep males in the sample in order to avoid complications of modeling labor supply decisions and to be as precise as possible about the type of prejudice we are estimating. We also restrict our sample to individuals between the ages of 18 and 65 who remain active in the labor market throughout their spell in the sample. We only consider individuals in two labor market states: unemployed or employed in a full-time job in the private sector.[18] Finally, we restrict the sample to individuals who at some point during the sample were employed at a manufacturing firm.

Following these restrictions, we are left with a sample of 114,984 males (8,714 blacks vs. 106,270 whites), of which 4,468 are unemployed (735 blacks vs. 3,733 whites). When defining the sample we face a trade-off between sample homogeneity and sample size. Because unemployed blacks represent a very small share of the population, in order to have a representative sample of these individuals (one that provides accurate estimates of moments related to job mobility) we need a large sample. Since the steady-state assumptions limit the sample size in the longitudinal dimension, we have to sacrifice the homogeneity of the sample. This is the main reason why we include individuals of all working ages (18-65 years-old) and education levels. Note that this does not conflict with our interpretation of the model. Age and education are strong predictors of individual skill, but the way skill is defined in the model captures all dimensions of skill, be them observed or not by the econometrician.

Workers' wages are weekly and measured at the time of the interview. Wages are observed at most twice per individual and are top-coded.[19] To deal with these issues we trim the top and bottom 2% of the wage distributions for black and white males. Using information on the number of working hours we convert wages to be hourly, which should further reduce the problem of top-coding.

Our source of firm-level data is the NBER-CES Manufacturing Industry Database.[20] To our knowledge, this is the only publicly available data set for the US with information on value-added that can be readily matched with CPS data. We match firm-level data to worker-level data by NAICS four-digit industry code. Our matched sample contains 76 distinct manufacturing industries from the year 2005. For each four-digit industry we compute the average value-added per worker per hour and interpret it as the level of production generated by the match between that worker-firm pair (i.e. $f(h, x)$). This means that, in our empirical application, technology differences across firms are only explained by the four-digit industry at which they operate.

[18]In addition, we excluded individuals with reported working hours outside the 35 - 70 hours interval.

[19]There is a literature about the misreporting of wages in the CPS. The main source of measurement error is believed to be over reporting at low levels and under reporting at high levels (see Bollinger (1998)).

[20]This database is a jointly produced by the National Bureau of Economic Research (NBER) and U.S. Census Bureau's Center for Economic Studies (CES). It contains annual industry-level data from 1958-2009 on output, employment and other variables for the 473 six-digit 1997 NAICS industries (see `http://www.nber.org/nberces`).

Finally, our sources of market-level data are the Job Openings and Labor Turnover Survey (JOLTS), the CPS and the Current Employment Statistics (CES), available from the Bureau of Labor Statistics webpage. To estimate labor market tightness in the same period in the manufacturing sector, we use the unadjusted series of job openings in the manufacturing sector from the JOLTS (JTU30000000JOL), the unadjusted unemployment rate series in the manufacturing sector from the CPS (LNU03032232) and the series of unadjusted employment in the manufacturing sector from the CES (CES3000000001).

4 Estimation

The relevant moments predicted by the model do not have a closed-form expression. Therefore, we make use of simulation methods to estimate the model's parameters. Specifically, we estimate the model by indirect inference (see Gouriéroux et al. (1993)). This procedure involves a simulated minimum distance estimator, in which some or all of the moments the procedure seeks to match are parameters from reduced-form models that capture important aspects of the 'true' data-generating process (i.e. the structural model). The parameters of these auxiliary models are a function of the structural parameters we seek to estimate.

The mechanics of an indirect inference procedure are the following. Let $\boldsymbol{\theta}$ denote the vector of structural parameters (to be specified in the next subsection), $\hat{m}^S(\boldsymbol{\theta})$ denotes the model-generated vector of parameters of the auxiliary models and $\hat{m}$ its empirical counterpart. (The contents of these vectors are defined in Section 4.2). The estimation procedure finds $\boldsymbol{\theta}$ such that the distance between the model-generated moments and their empirical counterparts is as small as possible, according to the following criterion function:

$$L_N(\boldsymbol{\theta}) = \frac{1}{2}\left(\hat{m} - \hat{m}^S(\boldsymbol{\theta})\right)' \Omega^{-1} \left(\hat{m} - \hat{m}^S(\boldsymbol{\theta})\right). \tag{16}$$

To obtain the theoretical moments we need to simulate the model. A description of the standard simulation algorithm of the model is provided in Appendix C. The weighting matrix Ω gives equal weight to all moments.[21] Standard errors are computed using a bootstrap procedure. We compute the moments for 200 resampled data sets and obtain a new set of parameter estimates based on each set of resampled moments.

[21] We also experimented with an *optimal* weighting matrix based on the variance-covariance matrix of the bootstrapped moments. The estimation results are not very different but, because transition rates in our data set are not precisely estimated compared to other moments, an *optimal* weighting matrix gives much more weight to wage moments. As a result, the fit of other moments (namely transitions from unemployment to job), which we consider equally important from an economic point of view, is quite poor using the *optimal* weighting matrix. In addition, when moment conditions are based on relatively few observations, as in this instance, it has been shown that equally-weighted distance matrices can perform better than *optimally*-weighted ones, see Altonji and Segal (1996).

In practice, to estimate the model we reformulate the minimization problem as a mathematical programming problem with equilibrium constraints (MPEC), as described by Su and Judd (2012).[22] The advantage of this method is to allow the program to search for the minimum without requiring that at each iteration all equilibrium conditions are exactly satisfied. In our case, rather than forcing the model to satisfy the free entry conditions at each iteration, we only require it to do so at the minimum. We augment the objective function $L_N(\boldsymbol{\theta})$ by the free entry conditions, imposing an arbitrarily small tolerance level for the restrictions to be satisfied and an arbitrary penalty parameter (akin to a Lagrange multiplier, but not optimally determined). The endogenous objects that are determined by these conditions $(x^{P\star}, G)$ are then included in the vector of exogenous parameters and chosen optimally to minimize the augmented objective function. In sum, this method allows for faster iterations, but at the cost of increasing the dimensionality of the problem (thereby requiring more iterations to find the minimum). This method proved somewhat more efficient in our case, but it need not be always the case (see Jørgensen (2013) for a thorough discussion). To minimize the augmented objective function we use a Nelder-Mead simplex algorithm. To find the global minimum we experimented with different initial values.

4.1 Econometric Specification

In order to make the model developed in Section 2 empirically operational we make certain parametric assumptions and calibrate some parameters. The population shares of worker types, m and $(1-m)$, are observed (92% and 8%, respectively). We specify the meeting function to be Cobb-Douglas with Constant Returns to Scale (CRS) and meeting elasticities equal to 0.5, i.e. $M(u_1 + u_2, v^P + v^N) = \lambda(u_1 + u_2)^{0.5}(v^P + v^N)^{0.5}$, where λ is the constant matching efficiency parameter to be estimated. The monthly discount rate, ρ, is set at 0.0043 (equivalent to 5% per annum).

We assume a multiplicative production function in the skill levels of workers and firms. When describing the model in Section 2 we assumed, for exposition purposes, that all skill distributions were the same, namely uniforms with support in the unit interval. We relax these assumptions in the empirical implementation of the model. We assume that all three distributions are log-normals, but allow the parameters of each distribution to differ. This allows us to summarize skill heterogeneity of this economy in six parameters, respectively the mean and standard deviation of white and black workers' skill distribution, $\{\mu_1, \mu_2, \sigma_1, \sigma_2\}$, and the mean and standard deviation of firms' technology, $\{\mu_x, \sigma_x\}$. To specify the log-normal distributions we transform the production function, replacing the rank of workers' and firms' skill by the inverse of the normal CDF, Φ^{-1}, which we then transform to log-normal distributions with mean, μ, and standard deviation, σ.

[22] We are grateful to Thomas Jørgensen for suggesting this method to us.

$$f_i(h,x) = \exp\{\mu_i + \sigma_i \Phi^{-1}(h)\} \exp\{\mu_x + \sigma_x \Phi^{-1}(x)\},\ i = 1,2.$$

Together these assumptions leave us with the following vector of parameters to estimate:

$$\boldsymbol{\theta} = \{\lambda, \delta, d, \pi, \mu_1, \sigma_1, \mu_2, \sigma_2, \mu_x, \sigma_x, \beta, b, \kappa\}$$

using data on individuals (indexed by i) at different months t, namely data on their race, $\bar{r}_i$, employment status, $\bar{s}_{i,t}$, log-wage, $\bar{w}_{i,j,t}$, log value-added per worker of the 4-digit manufacturing industry in which they are employed, $\nu_{j(i),t}$, (firms are indexed by j), and an estimate of the average market tightness in the manufacturing sector during the same time period ($\bar{\theta}$).

4.2 Auxiliary Models

The choice of auxiliary models is a key step in an indirect inference procedure. Our choices are based upon what we deem to be the crucial aspects of our theoretical model (in the sense that they capture aspects of the model that are informative about the value of certain parameters) and for which we have reliable empirical counterparts. Due to the complexity of the model, we can only provide an heuristic argument to support parameter identification.

4.2.1 Labor Market Transitions

Our model predicts three different rates of transition across labor market states. We summarize the information about transitions in the model in three moments: the average rate of transitions from job to unemployment for all employed workers and the average rates of transition from unemployment to job for white and black unemployed individuals. Their theoretical counterparts are given below:

$$jtu = 1 - e^{-\delta \times 1}, \tag{17}$$

$$utj_i = \int \left[1 - exp\left(-\lambda^W \sum_{j=\{P,N\}} \int \alpha_i^j(h,x) \frac{v^j(x)}{v^P + v^N} \, dx \times 1 \right) \right] \frac{u_i(h)}{u_i} dh. \tag{18}$$

The average rate of transitions from job to unemployment is essentially governed by δ, the parameter of the continuous Poisson process that determines when a match is destroyed. Similarly, the average rates of transitions from unemployment to job for white and black workers are informative about the Poisson process governing the meeting of workers and jobs. This can be seen clearly in equation (18), where $\lambda^W = \lambda(u_1 + u_2)^{0.5}(v^P + v^N)^{0.5}$. This expression also shows that the difference in mean unemployment-to-job transition rates across the two races is informative about employer prejudice (π and d). This is because these parameters affect differently the matching rates of workers of different races, via the matching sets, $\alpha_i^j(h,x)$ (see Section 7 for an illustration of these effects).

Under the steady state assumption, racial differences in unemployment rates conditional on skill are fully explained by differences in job finding rates across races. The assumption of a common exogenous job destruction rate across races is at odds with the data. In the US labor market black workers experience, on average, shorter employment spells compared to white workers. We agree with Lang and Lehmann (2012), who argue that mean differences in job-to-unemployment transition rates are likely to be quantitatively important to explain differences in unemployment rates across black and white individuals. Though appealing in practice, assuming exogenous differences in job destruction rates across races affects sorting patterns, via changes in the relative value of matches across types, and therefore differences in job finding rates. This imposes another source of discrimination which is not modeled and which also affects the complementarities between the race of a worker and a firm's output. Since our model does not generate endogenous differences in job destruction rates as result of employer discrimination and/or skill differences, we prefer to ignore this source of heterogeneity and quantify the sources of heterogeneity captured by the model.

4.2.2 Distribution of Wages and Value-added

The wages paid to workers of different skill and race are a key equilibrium object of the model. As can be clearly seen from the wage equation (equation (8)), the distribution of wages is informative about the distributions of workers' skill, the psychic cost of prejudice d and the technology distribution of firms. Wages also depend on reservation values and so they are also affected by the share of prejudiced firms (π) in the economy and the flow value of unemployment (b). To disentangle the skill distribution of black and white individuals, as well as employer prejudice parameters, we use moments from the wage distributions of black and white workers. In practice, we use the mean, standard deviation, skewness, kurtosis and the minimum of each of the two distributions. In the spirit of the argument in Flinn and Heckman (1982), the latter moment should be informative about the level of b.

To disentangle the contribution of workers' and firms' skill to the level of wages we make use of firm-level information. While not a direct counterpart to the distribution of firms' technology, the distribution of value-added per worker is informative about the level and variance of firms' technology.[23] We use the first two moments of the distribution of value-added.

[23] We interpret firms' value-added per worker as the counterpart of the match production $f(h,x)$. We do so because, when there is sorting on skill, firm value-added can no longer be taken as a direct measure of firms' technology.

4.2.3 Rent Shares Regressions

This auxiliary model draws on arguments made in the rent-sharing literature.[24] Conditional on the skill of workers and firms, h and x, the match production is divided up according to the worker and firm types and the bargaining parameter, β (see equation (8)). We estimate the following regression models separately for white and black employed individuals:

$$\bar{w}_{i,t} = \overline{\gamma}\nu_{j(i),t} + \epsilon_{i,t}, \tag{19}$$

where, $\epsilon_{i,t}$ is an i.i.d. disturbance and $\bar{w}_{i,t}$, $\nu_{j(i),t}$ respectively the log-wage and log-value-added.

The parameter $\bar{\gamma}$ describes the mean relationship between wages and firm value-added. These parameters of this regression on the population of white and black workers respectively are related to the structural parameters β, π and d in the following way. The level of the parameter $\bar{\gamma}$ is informative about rent-sharing in this economy, namely the value of β. Any differences in $\bar{\gamma}$ estimated separately in the population of black and white workers reveal information about the extent of discrimination and, therefore, of the magnitude of the two parameters that characterize employer prejudice (π and d). Indeed, if there were no prejudiced employers ($d = 0$), then, on average, white and black workers would get an equal share of value-added and $\bar{\gamma}$ estimated on the two populations of workers would be the same.

4.2.4 Labor Market Tightness

The ratio of vacant jobs to unemployed workers, $\bar{\theta} = \frac{V}{U}$, summarizes the degree of tightness in a labor market. The model counterpart of this statistic can be computed by integration of the accounting equations (14) and (15) respectively over the support of jobs' and workers' skill, which gives V and U. Given free entry of jobs and an aggregate meeting function, the number of firms in equilibrium, G, is fundamentally determined by the cost of posting a vacancy κ, via the two entry conditions (equations (5) and (6)). These entry conditions also determine the threshold technology level of the least efficient job operated by a prejudiced employer, $x^{P\star}$. In sum, given values for the other parameters of the model, market tightness informs the value of κ needed to satisfy the entry conditions.

5 Fit of the Model

In this section we evaluate the fit of the model. The first three columns of Table 1 report the set of simulated moments corresponding to three alternative model specifications. The first specification allows for both skill differences and discrimination resulting from employer

[24] Blanchflower et al. (1996) determine the share of profits accruing to workers in US manufacturing firms by matching CPS worker data with information on profits per sector (defined by 2-digit industry code), and regressing the former on the latter.

prejudice to govern workers' labor market outcomes. This is our preferred specification and is displayed in Column 1. We will also refer to this specification as the benchmark model. In the next specification (Column 2) we assume there are no prejudiced employers in the economy and therefore shut down the effect of hiring and wage discrimination on workers' labor market outcomes. The last specification we consider (Column 3) assumes black and white workers draw their skill levels from a common distribution, so that differences in labor market outcomes are the result of employer prejudice alone. The last two columns in Table 1 respectively present the empirical moments used in the estimation and their standard deviations.

Table 1: Goodness of Fit

		Model			Data	
		Unrestricted (1)	Skill (2)	Prejudice (3)	Mean (4)	Sd (5)
Transition Rates						
Unemployment to Job	White	0.306	0.302	0.308	0.313	0.011
	Black	0.258	0.272	0.246	0.263	0.025
Job to Unemployment		0.009	0.009	0.009	0.009	0.000
Wage Distribution						
Mean	White	2.895	2.906	2.859	2.993	0.005
	Black	2.652	2.715	2.706	2.679	0.015
Standard Deviation	White	0.490	0.501	0.466	0.485	0.004
	Black	0.445	0.459	0.499	0.444	0.009
Skewness	White	0.060	0.060	0.063	0.060	0.016
	Black	0.234	0.239	0.188	0.234	0.057
Kurtosis	White	2.152	2.149	2.153	2.240	0.016
	Black	2.212	2.233	2.105	2.343	0.075
Minimum Wage	White	1.918	1.910	1.932	1.946	0.003
	Black	1.813	1.851	1.788	1.714	0.050
Value-added Distribution						
Mean		4.227	4.250	4.315	4.288	0.006
Standard Deviation		0.603	0.608	0.610	0.548	0.007
Rent-share Regression						
Coefficient	White	0.678	0.679	0.677	0.671	0.002
	Black	0.641	0.658	0.651	0.596	0.007
Labor Market Tightness		0.385	0.394	0.361	0.381	0.013
Criterion		0.014	0.021	0.054	–	–

We first describe the empirical moments used to identify the parameters of the model. The top panel of Table 1 shows moments that govern worker mobility to and from unemployment. The first two moments are the average transition rates from unemployment to job of respectively white and black males in the sample. The third moment is the average transition rate from job to unemployment of all individuals in the sample. The observed difference in unemployment-to-job transition rates of whites and blacks is of 5 percentage points. The size of this difference is consistent with an average duration of unemployment for blacks that is 20 - 30% longer than that of whites, and which is reported in the various studies cited in Lang and Lehmann (2012). The job-to-unemployment transition rate common to all workers is 0.87%.[25]

The middle-top panel presents the first four moments of the distributions of log-wages of black and white workers. The differential in mean wages (in levels) of black and white workers is 31.4%, consistent with numbers found in the literature (see Lang and Lehmann (2012)) The standard deviations of both wage distributions have similar magnitudes, with whites' mean wage exhibiting slightly more variation. Both distributions are positively skewed, but the distribution of black wages is substantially more skewed to the right than that of whites' (0.234 vs. 0.060, that is four times as larger). The two distributions have very similar kurtoses and, as expected, the minimum wage (measured by percentile 2 of the wage distribution) of black workers is smaller than that of white workers (a 23.2% differential of wages in levels).

The middle-bottom panel displays the mean and standard deviation of the distribution of log-value-added of the 4-digit manufacturing industries in the sample. The next panel of Table 1 shows, for samples of white and black workers, the rent-share regression coefficients ($\overline{\gamma}$), which we also refer to as wage elasticity of firm value-added. With an elasticity of 0.671 against an elasticity of 0.596, white workers take, on average, a larger share of the match value compared to black workers. The last moment used in the estimation is the average labor market tightness in the manufacturing sector during the sample period.

Our preferred model (Column 1) fits the set of empirical moments quite well, although some moments are fitted better than others. Looking at the values across columns in the first rows of the top panel of Table 1 shows all models fit transition rates well. The benchmark model matches very well the difference in mean unemployment-to-job transition rates of whites and blacks (it generates a differential of 4.8% against a target differential of 5%). The fit of alternative models is worse. The model based on skill differences only generates a black-white transition rate differential of 3%, while for the model based on employer prejudice only that figure is 6.2%.

Regarding the fit of wage distributions moments, reported in the middle-top panel of Table 1, the benchmark model also performs better compared to alternative models. We focus on the first three moments of each distribution and emphasize how well each specification fits the differences in these moments across distributions. In the first column both predicted mean log-wages are below their empirical counterparts, but the difference is more substantial for whites. Therefore, the benchmark model understates the observed difference in mean log-wages across

[25] The difference in job-to-unemployment transition rates across races is 24% (1.02% and 0.82% for black and white employed workers respectively). The magnitude of this difference is in line with studies reported in the literature (see Lang and Lehmann (2012)).

races (25% vs. 31% in levels). The models in Columns 2 and 3 understate that differential even more (respectively 19% and 15%). The second moment of both log-wage distributions is fitted remarkably well by the models in Columns 1 and 2, while the model in Column 3 does clearly worse. A similar comment can be made about the performance of the three specifications in terms of fitting skewnesses.

We now turn our attention to the rent share regression coefficients, reported in the bottom panel of Table 1. The benchmark model captures the level of the coefficients well, but the difference in wage elasticities of value-added is half the magnitude present in the data (0.034 vs. 0.075). Both alternative models generate a smaller difference compared to the benchmark model (0.021 and 0.026 vs. 0.034).

In conclusion, although all specifications match quite well important subsets of moments, our preferred specification stands out as the one that comes closer to capturing simultaneously all the key aspects of the data: the difference in mean unemployment-to-job transition rates of whites and blacks, the black-white mean log-wage differential, the strong right-skewness of blacks' log-wage distribution and the difference in the wage elasticities of firm value-added. This information is summarized in the value of the minimization criterion at the bottom of Table 1, which is lower for the benchmark model (0.014) compared to the other two models (0.021 and 0.054).

6 Structural Parameter Estimates

Having argued that our model does a good job at matching the moments in the data, we now analyze its economic content. Table 2 reports the parameter estimates of the benchmark model, with standard errors reported in parentheses.

Table 2: Parameter Estimates

λ	δ	d	π	β	b	κ
0.826	0.009	8.322	0.490	0.051	4.500	537.861
(0.025)	(0.000)	(0.898)	(0.003)	(0.001)	(0.196)	(4.997)
μ_1	σ_1	μ_2	σ_2	μ_x	σ_x	
1.927	0.196	1.893	0.163	2.318	0.544	
(0.005)	(0.004)	(0.023)	(0.006)	(0.010)	(0.003)	

[1] Bootstrap standard errors based on 200 repetitions are reported in parenthesis.

The estimated job arrival and job destruction rates of the continuous-time Poisson processes are respectively 0.823 and 0.009. The size of these parameters is chiefly determined by the average transition rates between employment and unemployment in our sample. These moments are smaller than the average values reported in macro studies for the US economy over the same time period (see Shimer (2012)). This is consistent with the evidence on the relatively lower turnover in the manufacturing sector in the US (see Davis et al. (2013)).

The model produces evidence of employer prejudice against black workers. The loss of utility incurred by a prejudiced employer when matching with a black worker is estimated at 8.32 dollars per hour. This corresponds to 10.5% of the average productivity of a match involving white workers. The share of prejudiced employers is estimated at 49%. In the context of our model, these values provide a picture of discrimination in the US labor market as being driven by the presence of widespread and strong employer prejudice. This pattern is qualitatively similar to the one found in Bowlus and Eckstein (2002), who estimate that 56% of employers are prejudiced and that the disutility of hiring a black worker is 31% of white workers' productivity. However, this pattern is not consistent with the hypothesis advanced by Lang and Lehmann (2012). They '*(...) take the evidence from the surveys and the IAT* [Implicit Association Tests] *as suggesting that credible models of discrimination based on prejudice may rely on the presence of strong prejudice among a relatively small portion of the population and/or weak prejudice among a significant fraction of the population, but not on widespread strong prejudice*', Lang and Lehmann (2012), p. 969. They argue that strong and widespread employer prejudice is implausible, since it '*does not seem likely that a large proportion of employers (...) are willing to forego significant profits in order to avoid hiring blacks*', Lang and Lehmann (2012), p. 969. Quantitatively, our results come closer to a pattern of widespread and mild prejudice compared to other structural estimations based on the US labor market, but they are still far off. We offer a critical discussion of this hypothesis and our results in the concluding section of the paper.

Turning to differences in skill across black and white individuals, we estimate a mean black-white skill differential of 4.95%.[26] There are also differences, though smaller, in the standard deviations of the two distributions: whites' skill is more dispersed compared to blacks' (0.196 vs. 0.163). A simple Wald test of the equality of the two distributions (specified as a joint equality of the two parameters of wage distribution across races) is overwhelmingly rejected. There is obviously no available direct measure of workers skill to which we can compare our results. Nevertheless, these results are consistent with the persistent skill gap of black adults with respect to their white counterparts documented in detail in Neal (2006) — if anything his conclusions suggest the gap might be even larger. The worker bargaining parameter is estimated at 0.05, a value very close to zero. This estimate should be interpreted very cautiously, as the firm side data used in our empirical exercise is of limited quality.

Finally, the cost of posting a vacancy is estimated to be 537.861$ per hour. The order of magnitude of this parameter is surprisingly high. The first thing to note is κ and β are intrinsically linked. The high value of κ is in part driven by the enormous share of the surplus firms are able to keep $(1-\beta)$. Secondly, the estimate of κ is chiefly pinned down by one moment (market tightness), which is quite low in the US manufacturing sector.[27] We concede

[26] This ratio is calculated in the following way: $1 - exp(\mu_2 - \mu_1 - (\sigma_2^2 - \sigma_1^2)/2)$.

[27] Using data on the unemployment rate from the CPS, employment from the CES and on vacancies from the JOLTS, the estimated market tightness is 0.54, for the nonfarm economy, and 0.38, for the manufacturing sector.

that a low value of market tightness in the manufacturing sector is probably not related to very high vacancy posting costs. Our estimate of κ is likely to be picking up other aspects of the manufacturing sector not included in the model, like high capital costs.

7 Sorting on Race/Prejudice and Skill

Having estimated the parameters of the model we can simulate it and describe the main features of the equilibrium allocation. By doing so, a clear relationship emerges between matching decisions (captured by matching sets) and differences in labor market outcomes for workers of different races. To simulate the model we follow the methodology described in Appendix C, where the model is specified as detailed in Section 4.1. The first panel of Figure 1 displays the equilibrium strategies of all the agents in the economy (their matching sets). Figures 1a and 1b plot the contour lines of the set of matching sets of respectively prejudiced and nonprejudiced employers with white and black workers. Workers' skill (expressed in percentiles of the population skill distribution) is displayed on the horizontal axis and firms' technology (expressed in percentiles of the population skill distribution) on the vertical axis.[28] The solid lines indicate the bounds of the set of matching sets between white workers and prejudiced firms, while the dashed lines indicate the bounds of the set of matching sets between black workers and prejudiced firms. The region in the interior of each of the two pairs of lines contain all matches that produce a positive surplus.

A first observation is that the shapes of the two sets of matching sets embody sorting on skill: low(high)-technology firms match more often with low(high)-skill workers. Matches in the northwestern corner generate negative surpluses: high-technology firms have a high outside option that can only be compensated when they meet high-skill workers and production is high. Conversely, in the southeastern corner, the outside option of high-skill workers can only be outweighed when they match with high-technology firms.

A second observation concerns differences in whites' and blacks' sets of matching sets. Clearly, if there were no prejudice ($d = 0$) the two sets of matching sets would coincide. But, as can be seen by inspecting Figure 1a, the set of matching sets of blacks with prejudiced firms is contained in the one of whites. It implies that there is a range of technology levels for which prejudiced firms are willing to match with white workers, but not with equally skilled black workers. This is an instance of hiring discrimination stated in Corollary 2. The main mechanism underlying the difference in the two sets of matching sets is the direct effect of prejudice. It decreases the value of matches between black workers and prejudiced firms, to the point that even if when they are paid a lower wage compared to an equally able white worker (wage discrimination), some matches do not generate a positive surplus. This decrease in the set of matching

[28]Due to the estimated differences in skill distributions across races, there are no black individuals with extremely high levels of skill. This is why the dashed lines in all the four plots end before the solid lines.

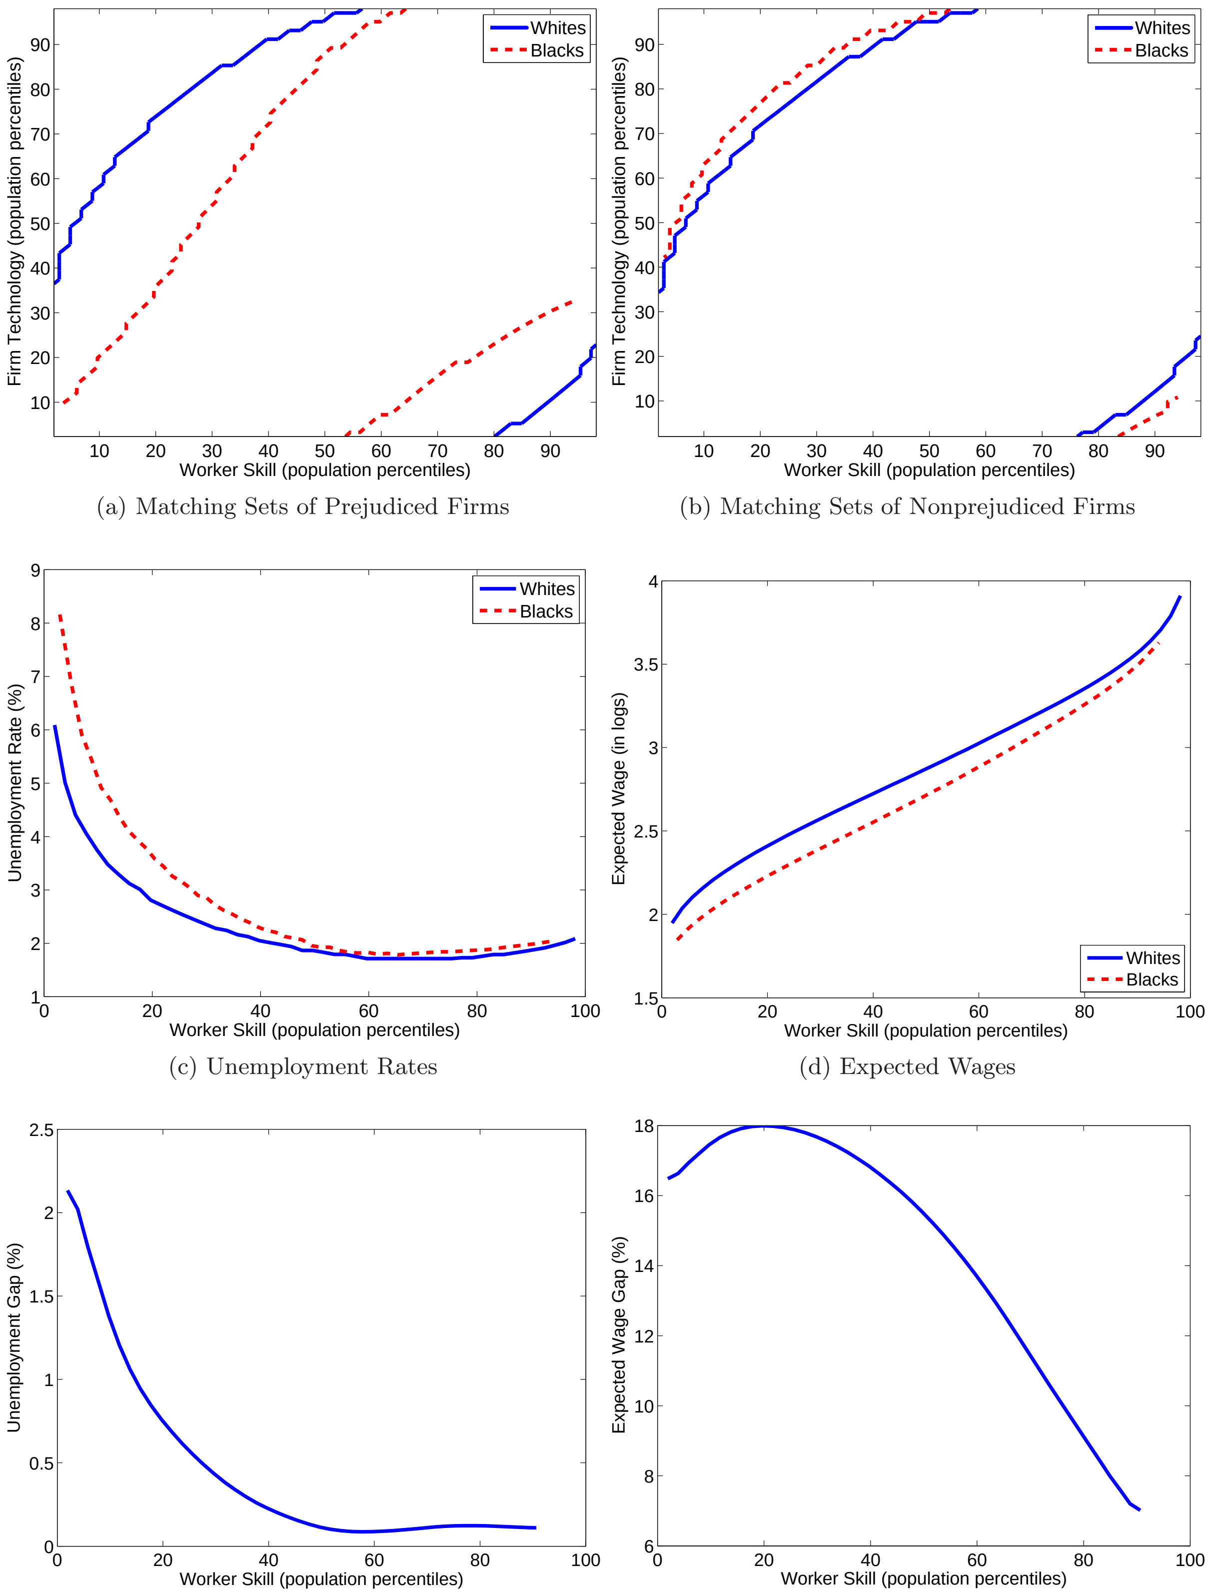

(a) Matching Sets of Prejudiced Firms

(b) Matching Sets of Nonprejudiced Firms

(c) Unemployment Rates

(d) Expected Wages

(e) Unemployment Rate Gap

(f) Expected Wage Gap

Figure 1: Simulation of Baseline Model

sets is asymmetric with respect to blacks' skill: it is stronger in the southwestern corner and milder in the northeastern corner. This follows from the way we modeled discrimination, where the psychic cost for a prejudiced employer upon matching with a black worker is independent of workers' and firms' skill levels. This implies low-skill blacks — who due to sorting match with low-technology firms — suffer relatively more from prejudice compared to their high-skill counterparts.

A different mechanism underlies reverse hiring discrimination (see Corollary 3). As can be seen in Figure 1b, the set of matching sets of black workers is larger than that of white workers in the northwestern and southeastern corners. There is an interval of technology levels for which nonprejudiced firms are willing to match with black workers but not with equally skilled white workers. The presence of prejudiced employers decreases blacks' opportunities in the labor market, which results in them having a lower outside option compared to equally skilled whites (see Proposition 1). This lower outside option makes matches of nonprejudiced employers of certain technology levels with blacks feasible compared to unfeasible matches with equally skilled whites.

The differences in labor market outcomes between black and white workers of similar skill depend on the relative magnitude of the various forms of discrimination present in the model, which in turn are determined by π and d and the share of whites and blacks in the population, m and $(1-m)$. In both Figures 1c and 1d the solid lines indicate white workers' outcomes and the dashed lines the outcomes of black workers, whereas the x-axis denotes the skill percentiles of the population of all workers. Figure 1c shows the unemployment rates of both types of workers for different levels of skill. More skilled workers experience lower unemployment rates. For both types of workers the distributions of unemployment rates across levels of skill describe a very nuanced U-shaped curve: unemployment rates decrease up until percentile 60 and increase slightly from there on. This results from the shape of matching sets and is a general feature of other search and matching models with sorting on skill (see Eeckhout and Kircher (2011)). A more important feature of this graph is that, whatever the level of skill of black workers, they experience higher unemployment rates compared to equally skilled white workers. This highlights what is already apparent in Figures 1a and 1b: that hiring discrimination dominates reverse hiring discrimination. More strikingly, the black-white unemployment rate differential is far higher for low levels of skill and becomes very small (approaching zero) as the level of skill increases. This can be seen more clearly in Figure 1e, which depicts the unemployment gap between black and white workers. This pattern illustrates the disproportionate effect of the psychic cost d on low-skill black workers' matching opportunities.

Finally, Figure 1d displays the differences in expected wages for blacks and whites with different levels of skill. The main patterns are similar to those in Figure 1c. Expected wages are increasing in workers' skill. For any level of skill, the expected wage of whites is always higher than that of equally skilled blacks. As can be seen more clearly in Figure 1f, the black-white expected wage differential is high for low-skill workers compared to high-skill workers. Similar to the effects of hiring discrimination, wage discrimination affects low-skill blacks proportionately more. In fact, part of the behavior of the wage gap along the skill distribution is driven by the decreasing differential between outside options of black and white workers, which is, in part, the result of the decreasing intensity of hiring discrimination. However, the main force at play is the decreasing intensity of direct wage discrimination ($-\beta d$ represents an ever decreasing proportion of wages, see equation (8)). The nonmonotonicity of the wage gap seems to be the result of the changing relative importance of two different forces. Across the first quintile of the population skill distribution, the wage gap increases, as the share of black workers employed in prejudiced firms in each skill percentile rises very fast, going from zero, at the infimum of the support, to a little over 25% by percentile 20. That is, in this region, the proportion of blacks employed in prejudiced firms is so small that, even though the intensity of wage discrimination is decreasing, the fact that the number of black workers who are exposed to direct wage discrimination is also increasing actually widens the wage gap. From percentile 20 onwards, the main force governing the behavior of the wage gap along the distribution of skill is the decreasing intensity of wage discrimination.

8 Decomposition of Wage and Employment Differentials

So far we have established that both skill differences and employer prejudice are important to match the patterns in the data pertaining to differences in blacks' and whites' labor market outcomes. We now assess their relative importance. We focus on the first, second (median) and third quartiles of the log-wage distribution and the mean unemployment rate. Columns 1 and 2 of Table 3 show the values of these moments, respectively for white and black employed individuals, generated by the parameter estimates of the benchmark model (Column 1 of Table 2). The next two columns of Table 3 show how much of the difference in moments can be explained by differences in skill and the existence of prejudiced employers, as a share of the same difference produced by both sources. Column 3 of Table 3 is obtained by simulating a model with the benchmark parameter estimates (column 1 of Table 2) except for the psychic cost, d, which is set to zero, to see how much of the gap can be attributed to skill differences alone. Analogously, Column 4 of Table 3 is obtained by simulating a model with the benchmark parameter estimates and setting the skill distributions of black and white workers to the mixture between the two skill distributions.

A number of remarks are in order. First, discrimination is quantitatively more important than skill differences for all the location measures of the wage distribution, with the exception of the third quartile. Second, prejudice is especially important to explain the differences in

Table 3: Decomposition of Wage and Unemployment Gaps

	Whites	Blacks	Share of gap (%): Skill	Prejudice
	(1)	(2)	(3)	(4)
Wage quartiles				
Q1	2.51	2.30	17.39	81.98
Median	2.89	2.61	44.30	58.42
Q3	3.27	2.99	57.04	38.60
Unemployment (%)	2.32	2.80	5.83	98.59

outcomes of black vs. white workers who earn lower wages (those who have lower levels of skill) and decreases monotonically as one moves towards higher wages. Again, this follows from the structure of the model, where the effects of prejudice on blacks' outcomes diminish as their level of skill increases (as discussed in greater detail in the previous section). Finally, employer prejudice explains all of the difference in average unemployment rates across races.

9 Counterfactual Analysis

In this section we conduct counterfactual analysis to evaluate the scope of alternative policy approaches in reducing differences in labor market outcomes of blacks with respect to whites. Beyond the impact of different policies on labor market outcomes, we are interested in quantifying their effects on individual and social welfare. We measure individuals' welfare by their expected present discounted lifetime utility value as described by the two equations below, respectively for workers and firms:

$$\mathrm{o}_i(h) = \frac{u_i(h)U_i(h) + \sum_{j=\{N,P\}} \int \gamma_i^j(h,x)W_i^j(h,x)\mathrm{d}x}{\ell_i(h)}, \tag{20}$$

$$\mathrm{o}^j(x) = \frac{v^j(x)V^j(x) + \sum_{i=\{1,2\}} \int \gamma_i^j(h,x)J_i^j(h,x)\mathrm{d}h}{g^j(x)}, \quad \forall x \,:\, V^j(x) \geq 0. \tag{21}$$

To evaluate the social desirability of different policies, we use a utilitarian social welfare function given by the sum of flow lifetime values. Equation (22) expresses the flow social welfare in the economy. It is composed of three elements: the flow output from existing matches net of the psychic cost of prejudice, the flow benefits received by unemployed workers and the flow recruitment costs incurred by vacant firms.[29]

[29] Alternatively, we could also consider a social welfare function that would not include the social losses due to prejudice. In our view, this would be a good choice if the social cost of prejudice were fully internalized by prejudiced employers. However, we think this choice is not adequate in our baseline model, where the utility cost of prejudice is partly transferred from prejudiced employers to black workers.

$$\rho \text{O} = b(u_1 + u_2) - \kappa(v^P + v^N) + \sum_{i=\{1,2\}} \sum_{j=\{N,P\}} \int\int \gamma_i^j(h,x)\left(f(h,x) - d\mathbf{1}_{\{i=2,j=P\}}\right) \mathrm{d}x\mathrm{d}h, \quad \forall x : V^j(x) \geq 0. \tag{22}$$

To characterize the welfare of different types of agents (black/white workers and prejudiced/nonprejudiced employers) in our economy we compute the total welfare of individuals belonging to said group, that is, the sum of flow lifetime values of agents who belong to a particular type (we name this measure *group welfare*).

Tables 4 and 5 report, respectively, the labor market outcomes and the levels of flow group and social welfare in four alternative scenarios. For each scenario Table 4 reports mean log-wages and unemployment rates of different types of workers, and mean vacancy rates and measures of active firms of different types of employers. Table 5 reports the level of group welfare for each type of agent, as well as the welfare gap across worker and employer types.[30] The last panel in Table 5 displays the level of social welfare (expressed in proportion of the level of social welfare in the baseline scenario), and the level of welfare accruing from each of its four components. In both tables, Columns 1, 2, 3 and 4 contain, respectively, the labor market outcomes and welfare measures in the baseline economy, in an economy in which the skill distributions of black and white workers are the same, in a scenario where a subsidy is given to employers who hire black workers and, last, in an environment where the psychic cost of prejudice is nontransferable. For each counterfactual we also describe the change in expected flow lifetime value for workers and employers of different types and skill levels with respect to the baseline economy, and report the respective sets of matching sets (see Figures 3, 4 and 5). In the following subsections we describe these four alternative scenarios in detail. We start by characterizing the baseline scenario, and then move on to characterize the different counterfactuals. For each of those we describe its rationale, as well as its effects on labor market outcomes and on social, group and individual welfare.

9.1 Baseline Economy

Figures 2a and 2b show the individual welfare of firms and workers of different types and skill levels in the baseline economy. The dashed lines denote individual welfare of prejudiced firms and black workers, whereas solid lines illustrate the welfare of nonprejudiced firms and white workers. Figures 2c and 2d display the welfare gap respectively across different types of firms and workers (prejudiced/nonprejudiced and black/white) with similar levels of skill. As we have documented in previous subsections, in the baseline economy black workers are, on average,

[30] Across workers' types the welfare measure used to compute this ratio is the mean welfare of each worker type. For firms this ratio is computed using the measure of group welfare, to take into account changes in the number of active firms of each employer type.

Table 4: Counterfactual Policy Simulations – Labor Market Outcomes

	Type	Policy: No Policy (1)	Equal Skill (2)	Employer Subsidy (3)	Nontransferable Prejudice (4)
Workers					
Mean Wage	White	2.896	2.899	2.889	2.905
	Black	2.652	2.760	2.835	2.757
Unemployment Rate (%)	White	2.332	2.320	2.322	2.320
	Black	2.809	2.778	2.395	3.146
	All	2.370	2.357	2.328	2.384
Employers					
Vacancy Rate (%)	Nonprejudiced	0.927	0.937	0.934	0.922
	Prejudiced	0.940	0.949	0.934	0.951
	All	0.933	0.943	0.934	0.936
Measure of Active Employers	Nonprejudiced	0.499	0.506	0.507	0.504
	Prejudiced	0.478	0.485	0.486	0.479
	All	0.977	0.991	0.993	0.982

[1] The subsidy paid to employers is 5$ per hour and the tax levied on employees is 0.4$ per hour.

Table 5: Counterfactual Policy Simulations – Group and Social Welfare

	Type	Policy: No Policy (1)	Equal Skill (2)	Employer Subsidy (3)	Nontransferable Prejudice (4)
Group Welfare	White	18.411	18.462	18.325	18.552
	Black	1.231	1.420	1.451	1.360
	Nonprejudiced	28.667	29.423	29.355	28.956
	Prejudiced	27.249	27.982	27.864	27.483
Welfare Gap (%)	Black/White	23.125	11.560	8.929	15.721
	Prejudiced/Nonprejudiced	5.012	4.892	4.873	5.685
Social Welfare	Total	1	1.019	1.016	1.007
	Production	80.40	81.95	81.72	81.19
	Unemployment utility	0.106	0.107	0.105	0.107
	Recruitment costs	4.902	5.028	4.986	4.946
	Utility costs of prejudice	0.243	0.253	0.278	0.200

[1] The subsidy paid to employers is 5$ per hour and the tax levied on employees is 0.4$ per hour.
[2] In Columns 1-4 social welfare (Total) is expressed as a proportion of the level of social welfare with respect to the level in the baseline economy.

unemployed more often and earn lower expected wages compared to white workers, and more skilled workers face lower unemployment rates and earn higher wages compared to their low-skill counterparts. When we use a more complete measure of individual welfare (one that combines information on wages and labor market dynamics) the same pattern persists. The welfare of workers is increasing in their level of skill. The pattern of the black-white welfare gap along the skill distribution follows that of the black-white wage gap, that is, it is larger for low-skill workers compared to high-skill workers.

Prejudiced firms experience higher vacancy rates and their share of the measure of active firms is lower than that of nonprejudiced firms. However, almost half of active employers are prejudiced. This is because the estimated share of prejudiced employers among the population of employers is 49% and the thereshold technology percentile $x^{P\star}$ is very close to zero. Similar to workers, the welfare of firms is increasing in their level of technology. Irrespective of their technology, the welfare of prejudiced employers is always lower than that of equally skilled nonprejudiced employers, but that difference is so marginal that is not visible to the naked eye in Figure 2a. This occurs because the share of black workers in the economy is very small and that part of the utility cost of prejudice can be transferred to black workers. Also note that the prejudice-nonprejudiced welfare gap decreases as the level of technology increases (cf. Figure 2c).

9.2 Premarket Affirmative Action

This counterfactual entails setting the skill distribution of black workers equal to that of whites. One can relate this exercise to policy interventions that intend to reduce the differences in skills between blacks and whites before they enter the labor market. We want to make the interpretation of skill in our model as general as possible. In our view, it can relate both to cognitive or noncognitive skills. As such, this counterfactual can capture a wide variety of real-world policy interventions aimed at improving the premarket skills of black individuals. In the US, policy programs along these lines have been implemented, aiming for instance at reducing existing social and urban segregation of black adolescents (see Lang and Lehmann (2012) and references therein).

The main effect of this counterfactual on labor market outcomes is an increase in the mean log-wage of black workers. The mean unemployment rate of black workers is but slightly reduced. This is consistent with the results obtained in the structural decomposition of mean log-wages and unemployment rates, which showed differences in unemployment rates across races were fully explained by employer prejudice. Put together, these changes lead to an increase in the group welfare of black workers. The direction of the effects on labor market outcomes and welfare of white workers is similar (albeit smaller). The mean log-wage of whites increases and their mean unemployment decreases, leading to an increase in the group welfare of white workers.

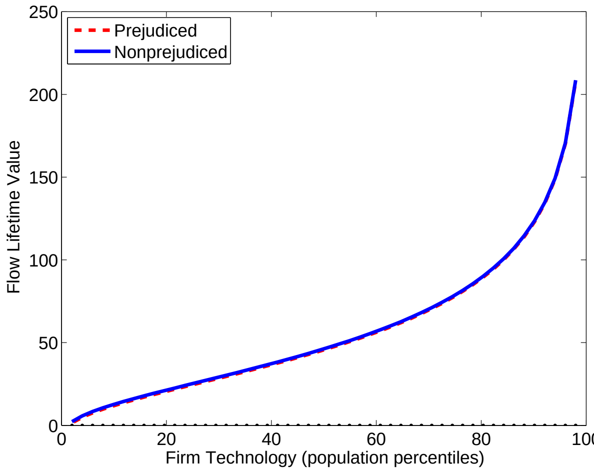

(a) Firms' Individual Welfare

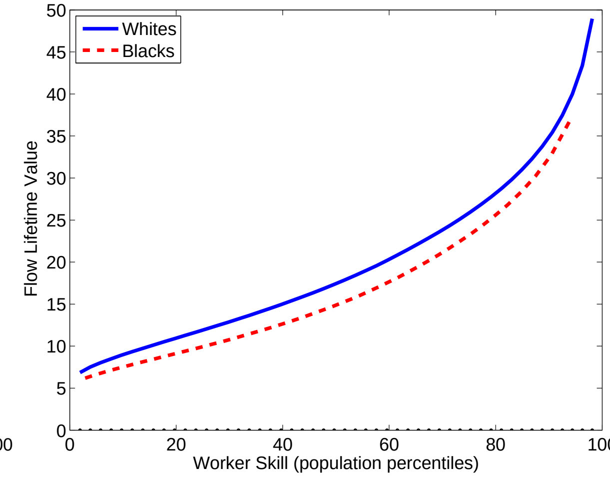

(b) Workers' Individual Welfare

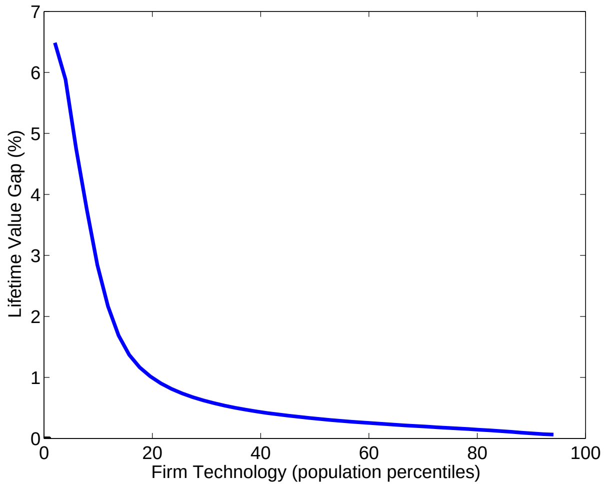

(c) Firms' Welfare Gap

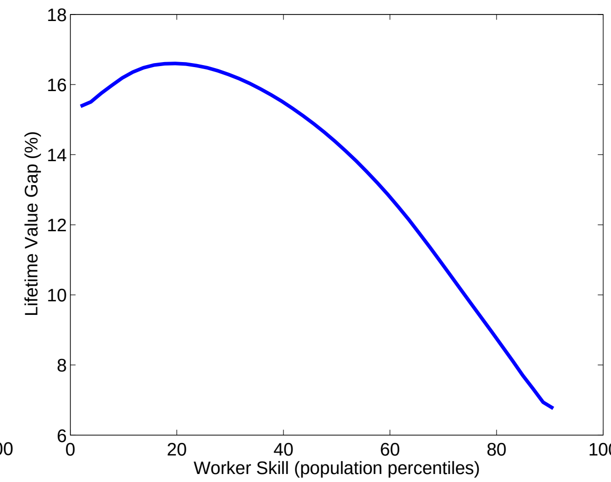

(d) Workers' Welfare Gap

Figure 2: Individual Welfare – Baseline economy

The sources of this increase are, however, different. While black workers receive an exogenous increase in their skill, all workers, irrespective of their race, benefit from the entry of more jobs, that reduces unemployment rates and pushes up workers' outside option and, by extension, their wages. This policy reduces the mean welfare gap of blacks with respect to whites by half its size in the baseline scenario (from 23.1% to 11.6%).

Inspection of Figure 3d shows the welfare gains are distributed quite asymmetrically across black workers.[31] Indeed, welfare increases proportionally more the higher the level of skill of black workers. A black worker who is born into the lowest percentiles of the skill distribution may even be worse off, as the higher standard deviation of the distribution of whites implies her level of skill may be lower than in the baseline scenario. The gradient of this increase in blacks'

[31] Note that the horizontal axis in Figure 3d denotes types' population percentiles. Hence, it is not suitable to draw comparisons between black and white workers. Because the majority of black workers see their level of skill increase in this counterfactual, a distributional assessment seems more informative if we make comparisons across scenarios looking at the change in welfare of workers born into the same skill percentile of their type distribution.

individual welfare along the distribution of skill is due to the exogenous changes in their skill distribution. This can be inferred by comparison with the change in welfare of white workers, which is quite homogenous (in proportional terms) along the skill distribution.

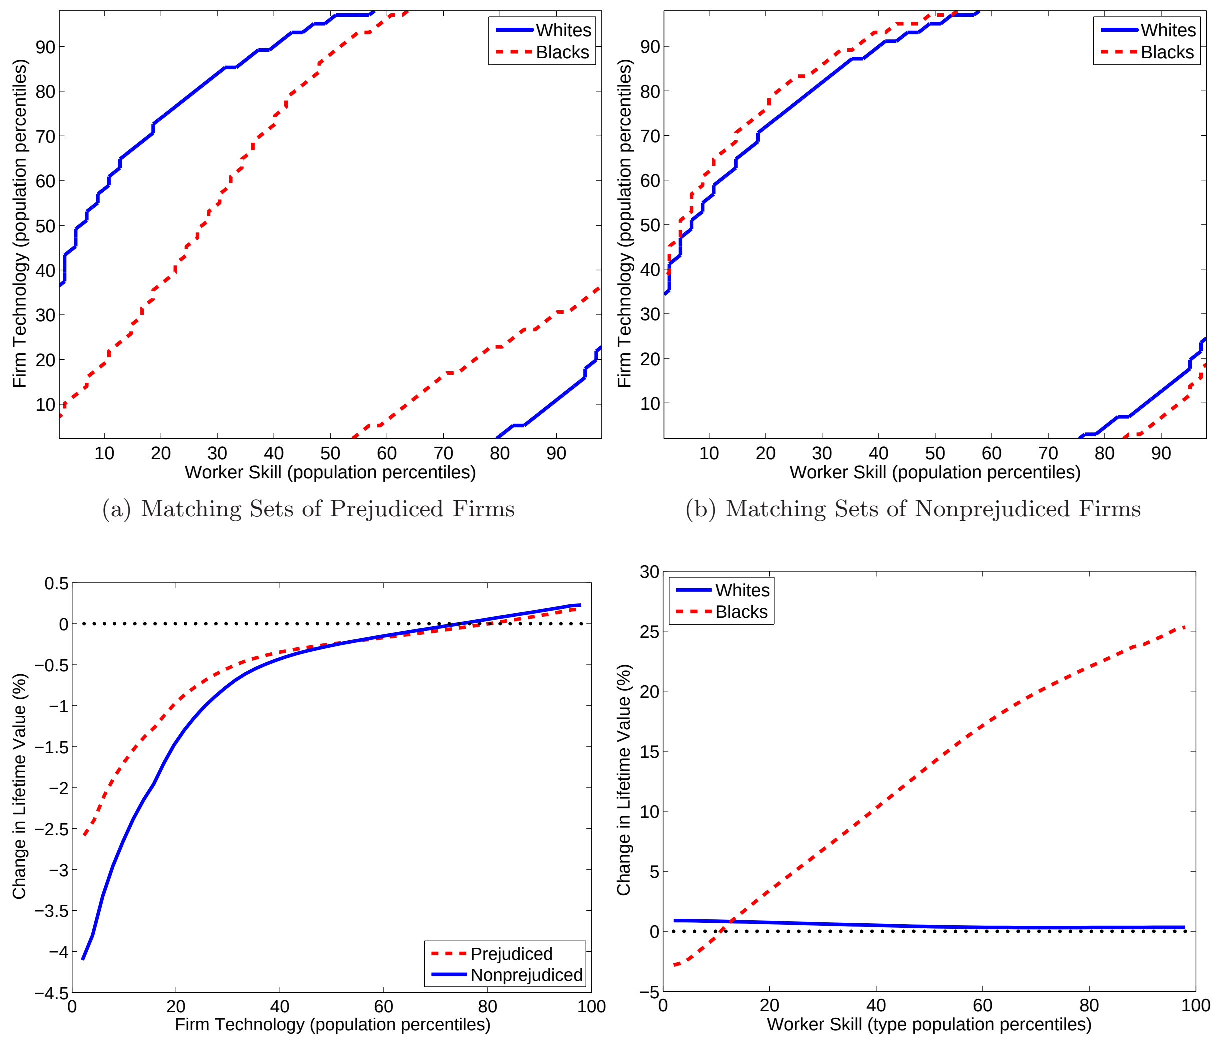

(a) Matching Sets of Prejudiced Firms

(b) Matching Sets of Nonprejudiced Firms

(c) Change in Firms' Individual Welfare

(d) Change in Workers' Individual Welfare

Figure 3: Impact on Individual Welfare of Premarket Affirmative Action

Group welfare of both prejudiced and nonprejudiced firms also increases, driven by a rise in the number of active employers of both types. Prejudiced firms as whole improve their position relative to nonprejudiced ones (the prejudiced/nonprejudiced welfare gap is reduced, though by a marginal amount). This said, because there are now more firms competing for the same number of workers, the individual welfare of most employers actually decreases. This is shown in Figure 3c. More than one cause seems to drive the patterns of distributional effects for firms. On the one hand, more efficient firms gain with this policy because the improvement in blacks' skill is concentrated on the top of the skill distribution and positive assortative matching ensures these gains fall more heavily on high-technology firms. On the other hand, increases in the measure of jobs are not homothetic: the lognormality assumption implies that increases in the mass of jobs are relatively higher at the bottom of the technology support, thereby increasing competition relatively more on this segment of the labor market. Since there are no perceptible changes in the sets of matching sets with respect to the baseline economy (compare Figures 3a and 3b with Figures 1a and 1b), the differential impact of this policy on prejudiced vs. nonprejudiced firms seems to result directly from the exogenous changes in the distribution of blacks' skill.

Not surprisingly, as a result of this policy social welfare increases. This rise is due to an increase in total production, which comes about through two different channels: an increase in the average quality of labor (black workers are more skilled) and a rise in the number of filled jobs. The change in production more than compensates the rise in social losses due to higher recruitment costs and utility costs due to prejudice. This policy does not affect the degree of skill and race/prejudice mismatch, since matching patterns are unaffected (cf. Figures 3a and 3b). Of course, one should bear in mind that this exercise does not take into account the social costs necessary to shift the distribution of black workers to that of whites.

9.3 Labor Market Affirmative Action

The second counterfactual involves paying a subsidy to any employer who hires a black worker. This exercise can be seen as an example of policy interventions that operate on the labor market indirectly, in the sense that the choice sets of agents are not directly constrained.[32] The constitutional legality of affirmative action policies is a contentious issue in the US. Opponents of such policies argue that treating two races differently is in direct violation to the Equal Protection Clause of the US Constitution. However, currently, the Supreme Court deems the practice constitutional (see Flabbi (2010) and references therein). These policies work to redress the disadvantage that black workers face compared to equally skilled white workers. This type of policy intervention has been put in place in several countries to push employment opportunities of disadvantaged workers, especially young individuals.

In our model economy, this subsidy has similar effects to those of the prejudice parameter d, but it works in the opposite direction (i.e. in favor of black workers) and it operates in matches involving either prejudiced or nonprejudiced employers. To make this exercise more meaningful we impose a balanced budget condition. The government's total outlays in subsidies to employers have to be financed by taxes levied on all employed workers. The impacts of this policy are clearly visible in the wage equation under this policy:

$$w_i^j(h,x) = \beta\left[f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} + \psi\mathbf{1}_{[i=2]} - \tau - \rho V^j(x)\right] + (1-\beta)\,\rho U_i(h)\,, \tag{23}$$

where τ is the tax level paid by employees and ψ the level of subsidy paid by the government to employers who hire a black worker. A more detailed specification of the equations correspond-

[32]The next policy counterfactual (Law Enforcement) provides an illustration of policy actions that restrict agents' choice sets.

ing to this counterfactual is given in Appendix D.1. The level of hourly payroll tax necessary to finance the subsidy can be expressed in the following way:

$$\tau = \frac{\psi\left((1-m)L - u_2\right)}{L - u_1 - u_2}. \tag{24}$$

A question that arises when simulating this policy is the level of subsidy to choose. We simulated the effects of this policy for different subsidy values. We choose to display the results from a realistic level of subsidy, 5$ per hour, and that results in an increase in social welfare.[33]

There are two direct effects of this policy: the mean log-wage of blacks increases because both the subsidy and the tax are transferable among match partners; and blacks' employment opportunities expand. As Figures 4a and 4b illustrate, both sets of matching sets with black workers expand, the one with prejudiced employers expanding the most. As a result of this, the mean unemployment rate of black workers decreases. This improves blacks' reservation values, leading to a positive feedback effect on wages. Both these changes lead to a substantial improvement in blacks' group welfare. Even in the presence of a sizeable skill gap, this policy is able to reduce an initial mean black-white welfare gap of 23% to a figure just below 9%.

As can be seen in Figure 4d, the proportional change in individual welfare is higher for low-skill black workers and decreases monotonically for high-skill workers. This is what we would expect based on our theoretical argument to justify decreasing wage and employment gaps along the distribution of skill. Figure 4d also shows a moderate decrease in lifetime values of white workers, which is concentrated on the low to middle range of the skill distribution. White workers' group welfare falls, and this loss in welfare is due to the fall in wages and in spite of a small decrease in the mean unemployment rate.

The effects on employers are similar to those of premarket market affirmative action, but now all employers (irrespective of their technology level and racial attitude) suffer a loss in expected lifetime utility. The main source of this change is the same: competition among employers increases, driving down their rents. Nonetheless, both employers' group welfare increases, due to the rise in the number of employers on the market. This policy affects negatively low-technology nonprejudiced employers proportionally more compared to their prejudiced counterparts. This seems to reflect the increase in the competition for labor on this segment of the market, due to a larger expansion of matching sets of prejudiced firms with black workers in this region (cf. Figures 4a and 3a). The monopsony power of nonprejudiced firms falls more in this region and more so than in equally skilled prejudiced firms. This policy leads to social gains and these occur via an increase in production, due to a higher number of filled jobs. Matches that were

[33]This subsidy level does not necessarily maximize social welfare. In fact, we have found values for the subsidy that yield a higher social welfare. However, increments in social welfare beyond a reasonable subsidy level are small. Our experiments also showed that the impacts of this policy on labor market outcomes are monotonic over a large discrete space of subsidy levels. More importantly, they indicate that increasing the level of subsidy beyond a certain level leads to ever greater social losses. Results are available upon request.

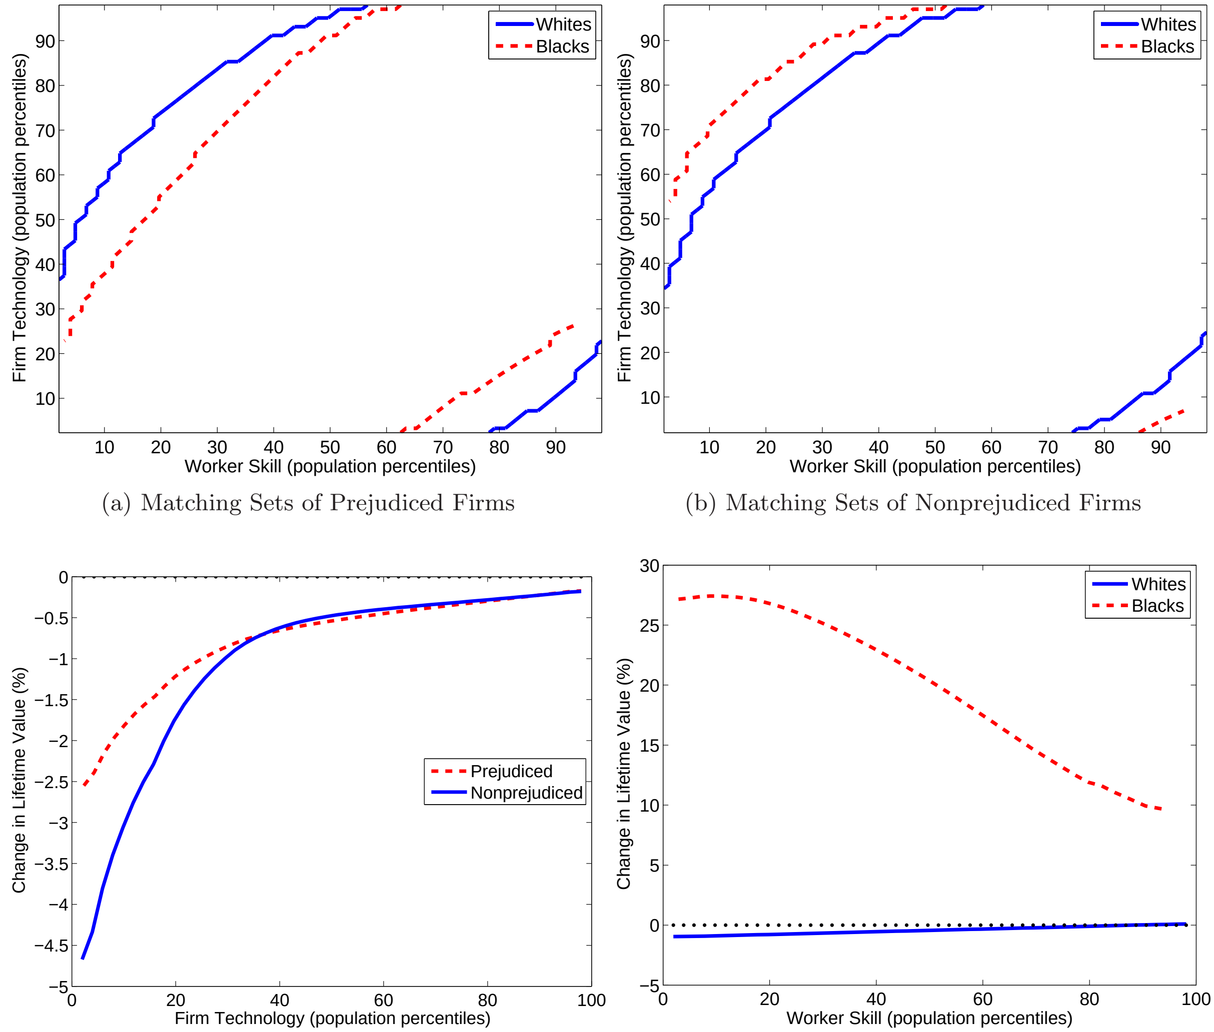

(a) Matching Sets of Prejudiced Firms

(b) Matching Sets of Nonprejudiced Firms

(c) Change in Firms' Individual Welfare

(d) Change in Workers' Individual Welfare

Figure 4: Impact on Individual Welfare of Labor Market Affirmative Action

previously not viable are now, on account of subsidies paid by more productive matches. These gains more than compensate the increase in recruitment costs and social losses due to prejudice.

9.4 Law Enforcement

The last counterfactual exercise we consider imposes the nontransferability of the utility cost of prejudice from prejudiced employers to black workers. One can view this exercise as being illustrative of the possible impacts of policy instruments that limit the contract space between these two types of workers and firms, in terms of reducing the ability of prejudiced employers to pay lower wages to equally skilled black workers. In other words, policy actions taken to enforce equal pay, that is, eliminate wage discrimination completely. Note, however, that in our modeling framework (where wages depend explicitly on outside options and search frictions are random), equal pay is unlikely to be fully enforced by making the utility cost of prejudice nontransferable.[34]

This kind of policy intervention seems a natural one to consider given that, in the US, any form of wage discrimination is illegal under the Equal Pay Act of 1963. In our baseline model, we assume that workers and employers are free to agree wages that violate equal pay. We believe this is a plausible description.[35] What we are implying with this characterization is that the existing law enforcement technology is not sufficiently effective to deter wage discrimination. Consequently, if the bargained wage is utility-increasing for both parties involved in the bargaining, it will be agreed upon and observed in our economy irrespective of it being illegal or not.[36]

In the baseline model wage discrimination against blacks occurs via two channels: the direct effect of d on wages and its effect through the outside option. When a black worker and a prejudiced firm negotiate the wage, the utility cost of prejudice enters the negotiation explicitly. Both parties recognize the presence of this cost and its decrease on the utility value of the match. They agree to share that loss according to their rent-sharing parameters. In this counterfactual we want to consider policies that prevent the first form of wage discrimination from occurring. We recognize it is difficult to connect this counterfactual to actual policies. Consistent with our previous description, we can think of it as the result of an improvement in the enforcement technology, allowing legal enforcement agencies to better observe at some feasible economic cost workers' skill and firms' technology and their reported outside options. In principle, a mix of monitoring and fines can be put in place to discourage the enactment of contracts involving a transfer of d from employers to workers via a lower wage.[37] Even if this story may seem somewhat far-fetched, we think it is still instructive to simulate the effects of an environment in which the ability of prejudiced employers to wage-discriminate is reduced.

To specify this counterfactual we change the match feasibility condition between black workers (type-2) and prejudiced firms (type-P) from equation (10) to equation (25),

$$\alpha_i^j(h,x) = \mathbf{1}\left\{(1-\beta)[f(h,x) - \rho U_i(h) - \rho V^j(x)] > d\mathbf{1}_{[(i,j)=(2,P)]}\right\}. \tag{25}$$

This condition states that a prejudiced employer will only match with a black worker if its flow surplus is greater than the psychic cost d. In addition to the match feasibility condition, we also change the payoffs of black workers and prejudiced employers when they are involved in matches with each other. Now prejudiced firms fully internalize the psychic cost of prejudice and

[34] In theory, this possibility can occur when this policy drives all prejudiced firms out of the market.

[35] To our knowledge, no evidence exists suggesting this law is fully enforced.

[36] In our view, to describe economic behavior, what matters are agents' incentives not legal constraints. If the economic incentives in place are such that contracts entailing wage discrimination are jointly privately efficient, we assume agents will take those opportunities.

[37] Whether the baseline model has more descriptive power compared to the model where d is nontransferable is an open question. Previous papers in the random-search-taste-based-discrimination tradition incorporate the former assumption. In principle, with better data we could test their validity.

so none of it is transferred to black workers' wages. See Appendix D.2 for a full mathematical description of this specification.

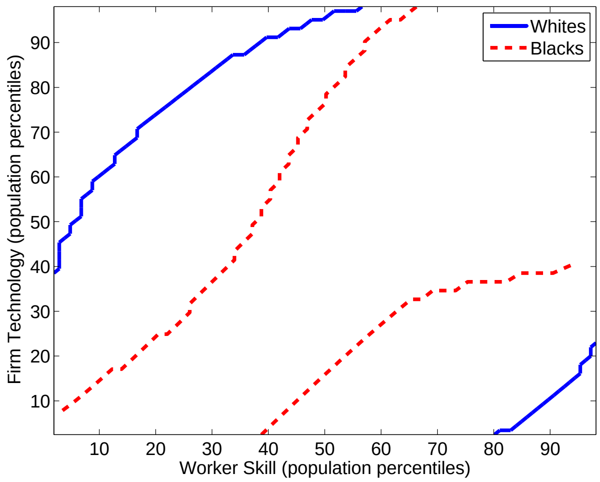

(a) Matching Sets of Prejudiced Firms

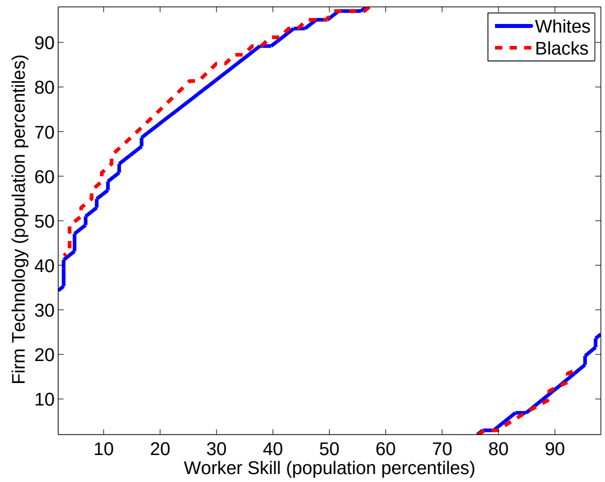

(b) Matching Sets of Nonprejudiced Firms

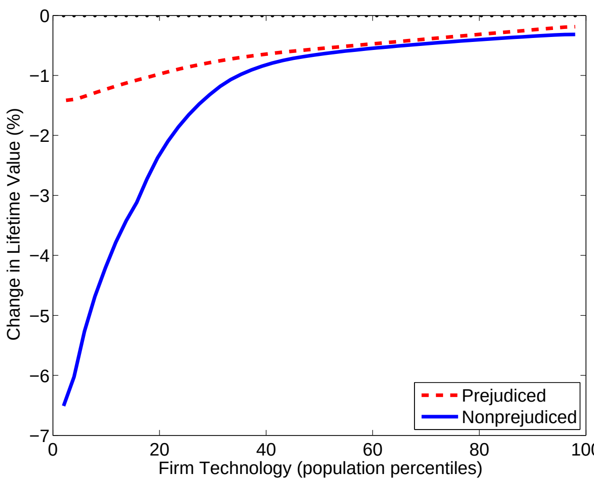

(c) Change in Firms' Individual Welfare

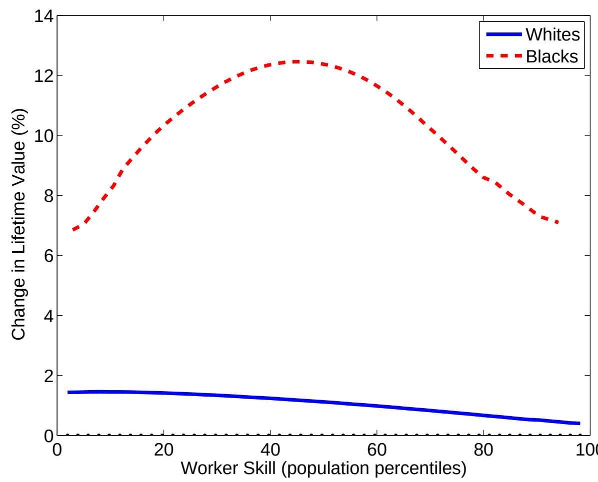

(d) Change in Workers' Individual Welfare

Figure 5: Impact on Individual Welfare of Law Enforcement

The fact that prejudiced firms now have to bear the full utility cost of prejudice implies two direct effects: certain matches between black workers and prejudiced firms are no longer viable, and the utility returns for prejudiced firms of being on the market are reduced. Despite the sizeable value of d, the small proportion of black workers on the market means that the utility loss experienced by prejudiced firms is small. Some low-technology prejudiced firms now have negative expected returns of entering the market and so decide to leave it ($x^{P\star}$ increases but very little). As these firms leave the market, opportunities arise for more nonprejudiced firms and more efficient prejudiced firms to join the labor market. Overall, the measure of active firms increases slightly, but this adjustment favors nonprejudiced employers the most, whose number on the market grows proportionally more compared to prejudiced employers (cf. the last column in the first panel of Table 4).

In terms of matching decisions, the main change occurs in the set of matching sets of prejudiced firms and black workers, which contracts towards the 45 degrees line (see Figure 5a). This contraction is asymmetric: it is deeper for low-technology firms compared to high-technology ones. Prejudiced firms now require a better match (that is, a skill level closer to the 45 degree line) with black workers in order to accept to match with them. This reduces skill mismatch in jobs involving black workers and prejudiced firms. Overall, the effect on group welfare of prejudice employers is positive, but of a very small magnitude. As can be seen in Figure 5c individual welfare decreases, and this loss affects low-technology firms the most.

The main beneficiaries of this policy are black workers. This increase is due to higher wages and in spite of a higher mean unemployment rate (cf. Column 4 in first panel of Table 4). Wages of black workers increase for two reasons: the extent of wage discrimination is reduced and the average match productivity of employed black workers improves. This effect follows from the contraction in the set of matching sets of black workers with both prejudiced and nonprejudiced employers. Indeed, Figure 5b shows that, under this policy, the degree of reverse hiring discrimination is substantially reduced due to an increase in black workers' reservation values vis-a-vis equally skilled white workers. Naturally, these changes in matching sets imply an increase in black workers' unemployment rate. As can be observed in Figure 5d, black workers in the middle of the skill distribution benefit proportionally more. This seems to be driven by the stronger contraction of the set of matching sets in this region, which pushes expected wages upwards due to an increase in average matching efficiency of matches involving this set of workers.

The effects on white workers' wages are similar but they are quantitatively smaller and originate from different sources (cf. Column 4 in first panel of Table 4). Wages are higher due to an outside option effect, resulting from increased competition for labor, as the number of active firms is higher. The mean white unemployment rate decreases slightly (the matching patterns of both types of employers with white workers are virtually unchanged). Finally, nonprejudiced employers' welfare increases as a group, reflecting a quantity effect. They benefit more than prejudiced employers and, as a result, the prejudiced-nonprejudiced welfare gap increases almost 2 percentage points. The individual welfare of nonprejudiced employers decreases, as they face greater competition, and, like in previous counterfactuals, low-technology nonprejudiced employers lose proportionally more compared to equally efficient prejudiced employers.

If one ignores the cost of developing and implementing the new enforcement technology, this policy brings about a small increase in social welfare. Total output of the economy increases, as does the total utility of unemployment, and social losses due to employer prejudice drop down. These gains outweigh the increase in recruiment costs. We draw one main lesson from this policy. In our baseline model racial prejudice can be seen as a negative externality that prejudiced employers impose on black workers. The optimal prescription in this case is to force agents responsible for the externality to fully internalize it. Then, only the most efficient ones will be able to afford it. As our counterfactual simulation shows, the combination of an adequate instrument to force prejudiced firms to internalize the cost of prejudice and free entry imply a substitution of low-skill prejudiced firms by higher-technology prejudiced firms, but especially by nonprejudiced firms of different technological levels. Not only does this reduce differences in labor outcomes of blacks with respect to whites, but more importantly, it increases the production level of the economy. Although in our view, this counterfactual has the most appealing rationale, its quantitative effects turn out to be very limited. This follows from the fact that the share of blacks is very small in this economy, and therefore the cost of prejudice does not significantly affect the market returns of prejudiced employers.

10 Conclusion and Discussion

In this paper we develop a search and matching model of the labor market with two-dimensional heterogeneous firms and workers to replicate key stylized facts pertaining to racial discrimination in the US labor market. The model successfully replicates facts regarding mean differences in wages and unemployment rates between white and black workers, as well as across their respective skill levels. We estimate the model using publicly available data for the US manufacturing sector. The model captures empirical moments retrieved from different data sources and pertaining to different aspects of the labor market.

We used the estimated model to make a number of contributions to the literature. We first characterized the equilibrium allocation of workers to jobs in a frictional environment in which workers differ in terms of a nonproductive attribute as well as skill, and where jobs have different technology levels and are operated by employers with different racial attitudes. The equilibrium allocation is characterized by positive assortative matching and wage and hiring discrimination, where these two channels operate with more intensity among low-skill workers compared to high-skill workers. Second, we undertook a structural decomposition to quantify the relative contribution of discrimination (resulting from employer prejudice) and skill differences across races to observed differences in employment and wages across blacks and whites. Finally, we performed conterfactual analysis to assess the scope of alternative policy approaches in improving labor market outcomes of black workers.

We end the paper with a critical discussion of our main findings. Our estimation results portray employer prejudice in the US manufacturing sector as being strong and widespread. Like Lang and Lehmann (2012), we are inclined to think that a lower degree or extent of employer prejudice is a more credible result. However, we want to emphasize that this inclination is not based on robust evidence, but on a prior belief about the strength of competitive forces operating on real-world markets. Influenced by Arrow's (1973) point that perfect competition should drive prejudiced firms out of the market, the literature sees the survival of a large fraction of strongly prejudiced firms as not credible. We are very cautious in taking our estimation results as definitive, because the data limitations we face are serious and our model does not capture all the relevant empirical moments (we come back to this point below), but we want to underline that the hypothesis formulated by Lang and Lehmann (2012) needs more empirical and theoretical backing.

Our empirical analysis shows that, given the structure of our model, this pattern of employer prejudice is important to fit the large differences in mean transition rates from unemployment to job across races. The law enforcement counterfactual exercise also showed that, even when the cost of prejudice is high, shutting down direct wage discrimination (so that discrimination only operates through the outside option and prejudiced firms are forced to fully internalize the cost of prejudice) in the presence of free entry has only a limited effect on driving prejudiced firms out of the market. We think this is due to the small share of black workers on the market, which means the losses of being prejudiced (in terms of expected returns of market activity) are also small. Our analysis also highlights that the ability of employers and workers to violate equal pay legislation is key for the survival of prejudiced employers. This, in turn, depend on the ability of anti-discrimination enforcement agencies to observe workers' skill and the productivity of jobs (occupations). More empirical evidence is needed to understand the quantitative importance of these mechanisms.

We believe there is ample room to make further progress by extending our modeling approach in other directions and, in particular, by trying to match other relevant empirical moments. In our model all matches are exogenously destroyed at the same rate. However, assuming there are no differences in job destruction rates is at odds with the data (and the literature emphasizes that these differences are large and important to explain racial unemployment gaps). It would be straightforward to impose different rates of exogenous job destruction across worker types, but we find this a rather ad hoc modeling strategy and that would confound the identification of employer prejudice. Alternatively, if the decision to fire a worker were also modeled, in principle it would imply black workers were fired more often, which would subsequently affect the decision to match with them in a negative way. We do not explore this conjecture further in this paper as doing so would require introducing idiosyncratic productivity shocks to worker-firm matches, adding significant complexity to the model and the estimation protocol. Another important limitation of our work is that we take the skill distribution of workers as given. In reality, skill acquisition is likely to be affected by the degree of labor market discrimination. It would be interesting to extend the model in this direction.

Finally, in our view a complete empirical evaluation of the model calls for a much richer data structure than what we had access to. In addition to producing more robust results, in principle, access to a matched employer-employee data set would allow us to estimate the degree of production complementarities. A milder functional form assumption on the production function would provide greater insight into the exact effect of differential sorting patterns governed by employer prejudice. A richer data set would also allow for stronger identification and for out-of-sample fit exercises to be conducted. For example, the patterns of worker segregation along education and race documented in Hellerstein and Neumark (2008) in the US labor market are consistent with the two forms of sorting present in our model.

References

Abowd, J., F. Kramarz and D. Margolis, "High wage workers and high wage firms," *Econometrica* 67 (1999), 251–333.

Altonji, J. G. and L. M. Segal, "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics* 14 (July 1996), 353–66.

Arrow, K., *The Theory of Discrimination*, volume 1 of *Discrimination in the Labor Markets* (Princeton, NJ: Princeton University Press, 1973), 3–33.

Bagger, J. and R. Lentz, "An Equilibrium Model of Wage Dispersion and Sorting," *mimeo* (2012).

Bartolucci, C., "Gender Wage Gaps Reconsidered: A Structural Approach Using Matched Employer-Employee Data," *Journal of Human Resources* (Forthcoming).

Becker, G. S., *The Economics of Discrimination* (University of Chicago Press, 1971).

———, "A theory of marriage: Part I," *The Journal of Political Economy* (1973), 813–846.

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

Black, D. A., "Discrimination in an Equilibrium Search Model," *Journal of Labor Economics* 13 (April 1995), 309–33.

Blanchflower, D., A. Oswald and P. Sanfey, "Wages, Profits and Rent-Sharing," *Quarterly Journal of Economics* 111 (1996), 227–251.

Bollinger, C., "Measurement Error in the CPS: A Nonparametric Look," *Journal of Labor Economics* 16 (July 1998), 576–594.

Borowczyk-Martins, D., G. Jolivet and F. Postel-Vinay, "Accounting for endogeneity in matching function estimation," *Review of Economic Dynamics* 16 (2013), 440–451.

Bowlus, A. J. and Z. Eckstein, "Discrimination and Skill Differences in an Equilibrium Search Model," *International Economic Review* 43 (November 2002), 1309–1345.

Cain, G. G., *Chapter 13 The economic analysis of labor market discrimination: A survey*, volume 1 of *Handbook of Labor Economics* (Elsevier, 1986), 693–785.

Carneiro, P., J. J. Heckman and D. V. Masterov, "Labor Market Discrimination and Racial Differences in Premarket Factors," *Journal of Law and Economics* 48 (2005), 1–39.

Charles, K. K. and J. Guryan, "Studying Discrimination: Fundamental Challenges and Recent Progress," *Annual Review of Economics* 3 (September 2011), 479–511.

Davis, S. J., R. J. Faberman and J. C. Haltiwanger, "The establishment-level behavior of vacancies and hiring," *The Quarterly Journal of Economics* 128 (2013), 581–622.

Eckstein, Z. and K. Wolpin, "Estimating the effect of racial discrimination on first job wage offers," *Review of Economics and Statistics* 81 (1999), 384–392.

Eeckhout, J. and P. Kircher, "Identifying Sorting — In Theory," *The Review of Economic Studies* 78 (2011), 872.

Flabbi, L., "Gender Discrimination Estimation In A Search Model With Matching And Bargaining," *International Economic Review* 51 (August 2010), 745–783.

Flinn, C. and J. Heckman, "New methods for analyzing structural models of labor force dynamics," *Journal of Econometrics* 18 (January 1982), 115–168.

Gouriéroux, C., A. Monfort and E. Renault, "Indirect inference," *Journal of applied Econometrics* 8 (1993), S85–S118.

Hagedorn, M., T. H. Law and I. Manovskii, "Identifying Equilibrium Models of Labor Market Sorting," Technical Report, National Bureau of Economic Research, 2012.

Hellerstein, J. K. and D. Neumark, "Workplace Segregation in the United States: Race, Ethnicity, and Skill," *Review of Economics and Statistics* 90 (August 2008), 459–77.

Jacquemet, N. and J.-M. Robin, "Assortative Matching and Search with Labour Supply and Home Production," *Centre for Microdata Methods and Practice, Institute for Fiscal Studies* (2012).

Jørgensen, T. H., "Structural estimation of continuous choice models: Evaluating the EGM and MPEC," *Economics Letters* 119 (2013), 287–290.

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

Lise, J., C. Meghir and J.-M. Robin, "Mismatch, Sorting and Wage Dynamics," *mimeo* (2013).

Lopes De Melo, R., "Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence," *mimeo* (2013).

Neal, D., *Why Has Black-White Skill Convergence Stopped?*, volume 1 of *Handbook of the Economics of Education*, chapter 9 (Elsevier, 2006), 511–576.

Neal, D. A. and W. R. Johnson, “The Role of Premarket Factors in Black-White Wage Differences,” *The Journal of Political Economy* 104 (1996), 869–895.

Petrongolo, B. and C. A. Pissarides, “Looking into the black box: A survey of the matching function,” *Journal of Economic literature* 39 (2001), 390–431.

Phelps, E. S., “The Statistical Theory of Racism and Sexism,” *The American Economic Review* 62 (1972), pp. 659–661.

Ritter, J. A. and L. J. Taylor, “Racial disparity in unemployment,” *The Review of Economics and Statistics* 93 (2011), 30–42.

Rosén, A., “Search, Bargaining, and Employer Discrimination,” *Journal of Labor Economics* 21 (October 2003), 807–830.

Shimer, R., “Reassessing the ins and outs of unemployment,” *Review of Economic Dynamics* 15 (2012), 127–148.

Shimer, R. and L. Smith, “Assortative matching and search,” *Econometrica* 68 (2000), 343–369.

Su, C.-L. and K. L. Judd, “Constrained optimization approaches to estimation of structural models,” *Econometrica* 80 (2012), 2213–2230.

Torres, S., P. Portugal, J. T. Addison and P. Guimarães, “The Sources of Wage Variation: A Three-way High-dimensional Fixed Effects Regression Model,” Technical Report, Discussion Paper Series, Forschungsinstitut zur Zukunft der Arbeit, 2013.

A Omitted proofs

In all the proofs presented in this section we assume the production function $f(h,x)$ satisfies the regularity conditions stated in assumption A0 in Shimer and Smith (2000) and strict supermodularity. Formally, we have that:

Assumption 1 (Production Function): *The production function $f(h,x)$ is nonnegative, symmetric, continuous, and twice differentiable, with uniformly bounded first partial derivatives on $[0,1]^2$. The production function $f(h,x)$ is strictly supermodular. That is, if $h > h'$ and $x > x'$, then $f(h,x) + f(h',x') > f(h,x') + f(h',x)$.*

Proof of Proposition 1. **Outside Option Effects**

Part 1 of Lemma 1 in Shimer and Smith (2000) states that, for any worker of any skill level, her unemployment value can only be smaller if evaluated at some alternative (non-optimal) matching set. This result applies in our environment as well.[38] In particular, it implies that, for any type-1 worker with skill h:

$$\rho U_1(h) \geq b + \lambda^W \sum_{j=P,N} \int \alpha_2^j(h,x) \left[W_1^j(h,x) - U_1(h) \right] \frac{v^j(x)}{v^P + v^N} \, \mathrm{d}x. \tag{26}$$

Subtracting $\rho U_2(h)$ to both sides of this inequality, substituting in the bargaining solution (equation (9)) and rearranging one obtains the following inequality:

$$U_1(h) - U_2(h) \geq \frac{d}{\rho} \times \frac{\frac{\lambda^W \beta}{\rho+\delta} \int \alpha_2^P(h,x) \frac{v^P(x)}{v^P+v^N} \, \mathrm{d}x}{1 + \frac{\lambda^W \beta}{\rho+\delta} \sum_{j=P,N} \int \alpha_2^j(h,x) \frac{v^j(x)}{v^P+v^N} \, \mathrm{d}x}. \tag{27}$$

If $\pi \in (0,1)$ then all workers face a positive probability of meeting a prejudiced firm due to random matching frictions, i.e. $\frac{v^P(x)}{v^P+v^N} > 0$, $\forall x$. Since we are characterizing equilibria in which at least some matches of every type are feasible, the integral in the numerator is always positive and so, when $d > 0$, $U_1(h) > U_2(h)$, $\forall h$.

Mutatis mutandis, one can prove that $V^N(x) > V^P(x)$, $\forall x$.

Proof of Corollary 1. **Wage Discrimination**

Take an arbitrary h and x. If $\pi \in (0,1)$ and $d > 0$, then $U_1(h) > U_2(h)$, $\forall h$ and :

$$\begin{aligned} w_2^P(h,x) =& \beta \left[f(h,x) - d - \rho V^P(x) \right] + (1-\beta) \rho U_2(h) < \\ & \beta \left[f(h,x) - \rho V^P(x) \right] + (1-\beta) \rho U_2(h) < \\ & \beta \left[f(h,x) - \rho V^P(x) \right] + (1-\beta) \rho U_1(h) = w_1^P(h,x). \end{aligned} \tag{28}$$

[38] The proof is available from the authors upon request.

We have proven that $w_1^P(h,x) > w_2^P(h,x)$, $\forall(h,x)$. *Mutatis mutandis*, one can prove that $w_1^N(h,x) > w_2^N(h,x)$, $\forall(h,x)$.

Proof of Corollary 2. **Type-2 Hiring Discrimination by Prejudiced Firms**

Using equation (9) one can write $S_1^P(h,x)$ as a function of $S_2^P(h,x)$:

$$S_1^P(h,x) = S_2^P(h,x) + \frac{d + \rho(U_2(h) - U_1(h))}{\rho + \delta}.$$

First, proving $d + \rho(U_2(h) - U_1(h)) > 0$ will imply $S_1^P(h,x) > S_2^P(h,x)$, $\forall(h,x)$. The proof is similar to the proof of Proposition 1 except that we are interested in the upper limit of the difference $U_1(h) - U_2(h)$. Specifically,

$$U_1(h) - U_2(h) \leq \frac{d}{\rho} \times \frac{\frac{\lambda^W \beta}{\rho+\delta} \int \alpha_1^P(h,x) \frac{v^P(x)}{v^P + v^N} \,\mathrm{d}x}{1 + \frac{\lambda^W \beta}{\rho+\delta} \sum_{j=P,N} \int \alpha_1^j(h,x) \frac{v^j(x)}{v^P+v^N} \,\mathrm{d}x}. \tag{29}$$

Using the same argument as in the proof of Proposition 1, we have $U_1(h) - U_2(h) < \frac{d}{\rho}$. Take an arbitrary (h,x) and assume that a type-2 worker will not be hired, i.e. $S_2^P(h,x) \leq 0$. She will suffer from hiring discrimination by prejudiced firms for all combinations of (h,x) such that $0 < S_1^P(h,x) < \frac{d+\rho(U_2(h)-U_1(h))}{\rho+\delta}$.

Recall that we are characterizing equilibria where some but not all matches of every kind are feasible, which implies $\exists(h,x) : S_1^P(h,x) > 0$ and $\exists(h',x') : S_1^P(h',x') \leq 0$. Part 2 of Lemma 1 in Shimer and Smith (2000) states that the values of unmatched agents are Lipschitz and thus continous. This result applies in our environment as well.[39] In particular, it implies that $S_i^j(h,x)$ is continous with respect to both h and x. Hence, by the intermediate value theorem, $\exists(h'',x'') : 0 < S_1^P(h'',x'') = \epsilon < \frac{d+\rho(U_2(h)-U_1(h))}{\rho+\delta}$.

We have shown that there exists at least one combination of (h'',x'') between a type-1 worker and a prejudiced firm such that $0 < S_1^P(h'',x'') < \frac{d+\rho(U_2(h)-U_1(h))}{\rho+\delta}$. Since the support of workers' skill is the same, $S_2^P(h'',x'')$ is well-defined and we know that $S_2^P(h'',x'') \leq 0$, i.e. a type-2 worker of the same skill level suffers hiring discrimination by the very same firm.

Proof of Corollary 3. **Type-1 Hiring Discrimination by Nonprejudiced Firms**

$S_2^N(h,x)$ can be expressed as a function of $S_1^N(h,x)$:

$$S_2^N(h,x) = S_1^N(h,x) + \frac{\rho(U_1(h) - U_2(h))}{\rho + \delta}.$$

Proposition 1 implies $S_2^N(h,x) > S_1^N(h,x)$, $\forall(h,x)$. The rest of the proof is analogous to the proof of Corollary 2.

Proof of Corollary 4. **Threshold Technology Differences**

[39] The proof is available from the authors upon request.

We first state a result that will be instrumental in showing both results of interest. It states that, for each type of firm, the value of posting a vacancy is a monotonically increasing function in firm's technology.

Proposition 2 (MONOTONICALLY INCREASING OUTSIDE OPTIONS):

(i) $V^j(x)$ *is monotonically increasing in* x*, for* $j = N, P$*; and*

(ii) $U_i(h)$ *is monotonically increasing in* h*, for* $i = 1, 2$*.*

The proof is, again, based on Part 1 of Lemma 1 in Shimer and Smith (2000). The value inequality lemma states that for an arbitrary x':

$$\rho V^j(x) \geq -\kappa + \lambda^F \sum_{i=1,2} \int \alpha_i^j(h, x') \left[J_i^j(h, x) - V^j(x) \right] \frac{u_i(h)}{u_1 + u_2} \mathrm{d}h. \tag{30}$$

Hence, for all $x_1 < x_2$

$$V^j(x_2) - V^j(x_1) \geq \frac{\frac{\lambda^F(1-\beta)}{\rho+\delta} \sum_{i=\{1,2\}} \int \alpha_i^j(h, x_1)\left(f(h, x_2) - f(h, x_1)\right) \frac{u_i(h)}{u_1+u_2} \mathrm{d}h}{\rho + \rho \frac{\lambda^F(1-\beta)}{\rho+\delta} \sum_{i=\{1,2\}} \int \alpha_i^j(h, x_1) \frac{u_i(h)}{u_1+u_2} \mathrm{d}h}. \tag{31}$$

Since we assume production complementarities, specified by a supermodular production function, $f(h, x_2) - f(h, x_1) > 0$ for all $x_1 < x_2$, which implies $V^j(x_2) - V^j(x_1) > 0$.

Mutatis mutandis, the proof is the same for $U_i(h)$.

We now turn to the proof of Corollary 4. In the type of equilibria we study we always assume, without loss of generality, that the least efficient nonprejudiced firm makes zero profit, i.e. $V^N(0) = 0$. Proposition 2 states $V^N(x)$ is monotonically increasing in x, which implies all nonprejudiced firms are active. Whenever $\pi \in (0, 1)$ and $d > 0$, by Proposition 1, $V^P(0) < 0$. Then, by Proposition 2 and continuity of $V^P(x)$, there exists a threshold technology level, $x^{P\star} > 0$, such that $V^P\left(x^{P\star}\right) = 0$.

B Equilibrium Definition

Definition 3 (EQUILIBRIUM): *Given exogenous parameters* L, m, d, π, ρ, β, b, κ, δ, *the production function* $f(h, x)$*, a meeting function* $M(u_1 + u_2, v^P + v^N)$ *and measures of firms and workers* $\ell_i(h)$*,* $g^j(x)$*, an equilibrium is a vector*

$$(\alpha_i^j(h, x), u_i(h), v^j(x), U_i(h), V^j(x), G, x^{P\star})$$

that solves the system of equations composed of the value functions of unmatched agents (equations (32) *and 33), the entry conditions (equations* (5) *and* (6)*), the measures of unmatched* *agents (equations (34) and (35)) and the matching indicator functions of all the agents participating in the economy $\alpha_i^j(h,x)$ (equation 10):*

$$U_i(h) = \frac{b + \frac{\lambda^W \beta}{\rho+\delta} \sum_{j=\{N,P\}} \int \alpha_i^j(h,x) \left[f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \rho V^j(x) \right] \frac{v^j(x)}{v^P+v^N} \, \mathrm{d}x}{\rho + \frac{\rho\lambda^W \beta}{\rho+\delta} \sum_{j=\{N,P\}} \int \alpha_i^j(h,x) \frac{v^j(x)}{v^P+v^N} \, \mathrm{d}x}, \tag{32}$$

$$V^j(x) = \frac{-\kappa + \frac{\lambda^F(1-\beta)}{\rho+\delta} \sum_{i=\{1,2\}} \int \alpha_i^j(h,x) \left[f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \rho U_i(h) \right] \frac{u_i(h)}{u_1+u_2} \, \mathrm{d}h}{\rho + \frac{\rho\lambda^F(1-\beta)}{\rho+\delta} \sum_{i=\{1,2\}} \int \alpha_i^j(h,x) \frac{u_i(h)}{u_1+u_2} \, \mathrm{d}h}, \tag{33}$$

$$u_i(h) = \frac{l_i(h)}{1 + \frac{\lambda^W}{\delta} \sum_{j=\{N,P\}} \int \alpha_i^j(h,x) \frac{v^j(x)}{v^P+v^N} \, \mathrm{d}x}, \tag{34}$$

and

$$v^j(x) = \frac{g^j(x)}{1 + \frac{\lambda^W}{\delta} \sum_{i=\{1,2\}} \int \alpha_i^j(h,x) \frac{u_i(h)}{u_1+u_2} \, \mathrm{d}h}, \; \forall x \, : \, V^j(x) \geq 0. \tag{35}$$

C Simulation Algorithm

The application of our chosen estimation method requires the equilibrium of the model to be solved numerically. This section describes how this is carried out in practice.

1. We set initial values for all the equilibrium objects (values, densities, matching sets, the total number of firms in the economy and the threshold technology level of prejudiced firms) and parameters of the model $\boldsymbol{\theta} = \{\lambda, \delta, d, \pi, \mu_1, \sigma_1, \mu_2, \sigma_2, \mu_x, \sigma_x, \beta, b, \kappa\}$. We discretize the supports and measures of the skill distributions of workers and firms in the economy.

2. Given all initial values we find new values for G and $x^{P\star}$ such that $V^P\left(x^{P\star}\right) = 0$ and $V^N(0) = 0$

3. Using the initial values and updated G and $x^{P\star}$ we iterate over equations (32) and (33) to determine $U_i(h)$ and $V^j(x)$, at each stage updating the region of feasible matches determined by equation (10).

4. New values for $u_i(h)$ and $v^j(x)$ are obtained from equations (14) and (15).

5. Given the values determined in the previous step, new values of u_i and v^j as well as λ^W are determined.

6. Steps 3 through 5 are updated until the endogenous distributions $u_i(h)$ and $v^j(x)$ converge.

D Counterfactual Specifications

D.1 Labor Market Affirmative Action

To specify this counterfactual we need to alter the value of employment for workers and of a filled job for firms. These will, in turn, change the match feasibility conditions and the wage equation of the model. The value of being matched respectively for a worker and an employer are given by the following equations:

$$\rho W_i^j(h,x) = w_i^j(h,x) - \tau + \delta\left[U_i(h) - W_i^j(h,x)\right], \tag{36}$$

$$\rho J_i^j(h,x) = f(h,x) + \psi \mathbf{1}_{[i=2]} - d\mathbf{1}_{[(i,j)=(2,P)]} - w_i^j(h,x) + \delta\left[V^j(x) - J_i^j(h,x)\right] \tag{37}$$

where τ is the tax level paid by employees and ψ the level of subsidy paid by the government to employers who hire a black worker. The implied match feasibility conditions can be written in the following way:

$$\alpha_i^j(h,x) = \mathbf{1}\{f(h,x) - d\mathbf{1}_{[(i,j)=(2,P)]} - \tau + \psi\mathbf{1}_{[i=2]} - \rho U_i(h) - \rho V^j(x) > 0\}. \tag{38}$$

D.2 Law Enforcement

In a model with a nontransferable cost of prejudice, the expressions for the value functions remain unaltered and are given by equations (1), (2), (3) and (4). However, as the cost of prejudiced is no longer transferable from the employer to the worker, it is not present in the match surplus expression or in the wage equation, which are now written thus:

$$S_i^j(h,x) = \frac{f(h,x) - \rho U_i(h) - \rho V^j(x)}{\rho+\delta}, \tag{39}$$

$$w_i^j(h,x) = \beta\left[f(h,x) - \rho V^j(x)\right] + (1-\beta)\rho U_i(h). \tag{40}$$

Because prejudice employers have to incur the total cost of prejudice when they match with a black worker, a prejudiced employer will only match with a black worker if his flow surplus in that match is greater than the psychic cost d:[40]

$$\alpha_i^j(h,x) = \mathbf{1}\left\{(1-\beta)[f(h,x) - \rho U_i(h) - \rho V^j(x)] > d\mathbf{1}_{[(i,j)=(2,P)]}\right\}. \tag{41}$$

Finally, in equilibrium, the reservation values of prejudiced employers and black workers are given by the following equations:

$$V^P(x) = \frac{-\kappa + \frac{\lambda^F(1-\beta)}{\rho+\delta}\sum_{i=\{1,2\}}\int \alpha_i^j(h,x)\left[f(h,x) - \frac{d}{1-\beta}\mathbf{1}_{[i=2]} - \rho U_i(h)\right]\frac{u_i(h)}{u_1+u_2}\,\mathrm{d}h}{\rho + \frac{\rho\lambda^F(1-\beta)}{\rho+\delta}\sum_{i=\{1,2\}}\int \alpha_i^j(h,x)\frac{u_i(h)}{u_1+u_2}\,\mathrm{d}h}, \tag{42}$$

$$U_2(h) = \frac{b + \frac{\lambda^W\beta}{\rho+\delta}\sum_{j=\{N,P\}}\int \alpha_2^j(h,x)\left[f(h,x) - \rho V^j(x)\right]\frac{v^j(x)}{v^P+v^N}\,\mathrm{d}x}{\rho + \frac{\rho\lambda^W\beta}{\rho+\delta}\sum_{j=\{N,P\}}\int \alpha_2^j(h,x)\frac{v^j(x)}{v^P+v^N}\,\mathrm{d}x}. \tag{43}$$

[40] Note that this is the relevant match feasibility condition since $(1-\beta)S_2^P(h,x) > d \Rightarrow S_2^P(h,x) > d/(1-\beta) > 0$. for $d > 0$.