

NBER WORKING PAPER SERIES

PREJUDICE AND THE ECONOMICS OF DISCRIMINATION

Kerwin Kofi Charles
Jonathan Guryan

Working Paper 13661
<http://www.nber.org/papers/w13661>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2007

The authors thank Gary Becker, John Bound, Charlie Brown, Henry Farber, Larry Katz, Kevin Lang, Kevin Murphy and Paul Oyer for helpful discussions regarding this work, seminar participants at the University of Chicago, Harvard, Wharton, CUNY, Yale, Brown, the Institute for Research on Poverty, Princeton and the University of Georgia for their suggestions and Phil Wengerd for research assistance. Guryan thanks the University of Chicago GSB and the Industrial Relations and Education Research Sections at Princeton University for research support. This research was also funded in part by the George J. Stigler Center for the Study of the Economy and the State at the University of Chicago Graduate School of Business. All errors are our own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Kerwin Kofi Charles and Jonathan Guryan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Prejudice and The Economics of Discrimination
Kerwin Kofi Charles and Jonathan Guryan
NBER Working Paper No. 13661
December 2007
JEL No. J01,J7

ABSTRACT

This paper tests the predictions about the relationship between racial prejudice and racial wage gaps from Becker's (1957) seminal work on employer discrimination - something which has not previously been done in the large economics discrimination literature. Using rich data on racial prejudice from the General Social Survey, we find strong support for all of the key predictions from Becker about the relationship between prejudice and racial wage gaps. In particular, we show that, relative to white wages, black wages: (a) vary negatively with a measure of the prejudice of the "marginal" white in a state; (b) vary negatively with the prejudice in the lower tail of the prejudice distribution, but are unaffected by the prejudice of the most prejudiced persons in a state; and (c) vary negatively with the fraction of a state that is black. We show that these results are robust to a variety of extensions, including directly controlling for racial skill quality differences and instrumental variables estimates. We present some initial evidence to show that racial wage gaps are larger the more racially integrated is a state's workforce, also as Becker's model predicts. The paper also briefly discusses familiar criticisms and extensions of the standard Becker model, including an argument of our own which, like some recent work, shows that the model's main predictions can be shown theoretically to survive the effects of long run competition.

Kerwin Kofi Charles
Harris School of Public Policy
University of Chicago
1155 East 60th Street
Chicago, IL. 60637
and NBER
kcharles@uchicago.edu

Jonathan Guryan
University of Chicago GSB
5807 S. Woodlawn Ave.
Chicago, IL 60637
guryanj@nber.org

1. Introduction

Becker's (1957) seminal "The Economics of Discrimination" launched the formal analysis of labor market discrimination among economists. Becker's analysis focused on the relationship between racial prejudice among whites and discrimination against racial minorities in a competitive model. In contrast to much of the contemporaneous literature, Becker formalized the definition of racial preferences, depicting them as an aversion to cross-racial interaction. In a series of models, he analyzed the effect of the possession of such preferences among customers, co-workers and employers on black relative wages. Since the publication of Becker's work, discrimination has been one of the most intensely studied topics in economics; theorists have posited different explanations for racial wage gaps including many which do not depend on racial animus, and there have been hundreds of empirical studies devoted to measuring and attempting to explain wage differences.¹ Yet, curiously, in the massive and growing discrimination literature, there is no paper of which we are aware that tests the sharp, yet subtle predictions of Becker's original prejudice model about the relationship between racial wage gaps and prejudice.² In this paper, we attempt to address this omission.

In the short-run version of Becker's employer discrimination model, racial prejudice causes some employers to regard black workers as more expensive than they truly are. Market pressures cause blacks to be hired by the least prejudiced employers in the market and to sort away from those with the highest levels of prejudice. In equilibrium, Becker shows that black relative wages are determined by the most prejudiced employer with whom they come into contact – the marginal discriminator. Racial wage gaps, in Becker's formulation, are determined by the prejudice of this marginal person, and not by the average level of prejudice among all employers. More generally, given the sorting mechanism described above, and since blacks constitute a relatively small share of the labor force in almost every labor market, wage gaps will generally be determined by variation in the level of prejudice of those in the lower tail of the prejudice distribution; how prejudiced the most prejudiced employers are should not matter at all for wages in Becker's

¹ Some key explanations for discrimination that do not depend on racial prejudice include imperfect information in statistical discrimination models (Aigner and Cain (1977), Altonji and Pierret (2001)), imperfect competition in dual labor market and local monopsony models (Doeringer and Piore (1971), Black (1995)), and racial difference in productivity (Neal and Johnson (1996)).

² Although there has been very little work in economics studying racial prejudice directly, two exceptions are Cutler, Glaeser and Vigdor (1999) and Card, Mas, and Rothstein (2006). Both of these papers study the relationship between prejudice and residential racial segregation, but do not examine the wage relationships that are our focus.

framework. Finally, since the ease with which blacks are sorted to less prejudiced firms varies inversely with the number of blacks in the market, holding the level of prejudice constant, equilibrium wages for blacks should vary negatively with the number of blacks in the market.

Despite its sharp predictions, one possible reason that Becker's model's predictions about the connection between prejudice and wages have not been empirically tested is Arrow's (1972) famous criticism. Arrow argues that since prejudiced employers sacrifice profits by discriminating, such employers are ultimately driven from the market in the long run in a competitive setting. Indeed, Arrow memorably remarked that the employer discrimination model "predicts the absence of the phenomenon it was designed to explain."³ However, casual empiricism suggests that prejudicial feelings of the most odious sort were a feature of the American landscape for many scores of years, and logic dictates that these views have had *something* to do with racial discrimination against minorities. Supporting the view that prejudice may matter for wage determination, consider Figure 1 in which we relate the black-white wage gap in each of the nine Census divisions to the mean response among whites in those divisions to two questions from the General Social Survey about racial prejudice. The figure shows that blacks have much lower relative wages in places where a higher fraction of whites report opposing interracial marriage, or where whites would not vote for a black President. These figures do not confirm that a mechanism like that outlined by Becker links prejudice and wages, but they are suggestive of such a relationship.⁴

Arrow himself and many subsequent authors have shown that the connection between prejudice and wages posited by Becker can survive in the long run, if the market is not perfectly competitive, if there are important non-convexities, or if there are informational problems. Before performing our empirical tests of the Becker employer prejudice model, we briefly review this literature. In addition, we outline an argument of our own which shows that, indeed, the predictions of the original Becker model can be shown to hold in the long run, even if the market is fully competitive. We argue that most previous treatments of the Becker model make an

³ Arrow (1972), p. 192.

⁴Further potential evidence about the possible importance of racial prejudice may be found in a widely-known recent study by Bertrand and Mullanaithan (2004), who find that resumes with black sounding names sent to potential employers received fewer call backs than did other resumes. Although statistical discrimination might account for these results as suggested by later results on naming conventions studied by Fryer and Levitt (2004), the fact that so many personal traits of likely interest to employers are explicitly controlled for on the false resumes leaves open the possibility that some other force, like racial prejudice, might be at work.

implicit assumption that an individual's racial prejudice is not carried with the person across the different labor market roles he plays. We argue that relaxing this peculiar implicit assumption leads to conclusions about the equilibrium relationship between prejudice and wage gaps not markedly different from those derived in Becker's original formulation.

To empirically test the Becker model, we use the rich information on racial sentiments available in multiple waves of data from the General Social Survey (GSS). We summarize this information on racial prejudice, describing its levels among different sub-groups of whites and its evolution over time. In a series of analyses, we relate prejudice among whites in a state to the observed racial wage gap, using wage data from the Current Population Survey (CPS). The data provide overwhelming support for each of the predictions of Becker's model. First, we find that racial wage gaps are much more closely related to the level of prejudice of the "marginal" person in the distribution than they are to average levels of prejudice. We further show that it is only prejudice in the left tail of the prejudice distribution that seems to matter for wage gaps; wages do not vary at all with the prejudice of the most prejudiced persons in a state. Importantly, the foregoing results are from regressions that control for the racial makeup of states. Finally, we show that the fraction of a labor market that is black, holding prejudice in the state constant is strongly negatively related to racial wage gaps, just as the prejudice model predicts. These results are robust to a variety of robustness tests and extensions. On the whole, we believe our results are broadly suggestive of an important role for racial prejudice among whites in explaining differential labor market outcomes by race.

The remainder of the paper proceeds as follows. The next section briefly reviews Becker's employer racial distaste model, and reviews criticisms and extensions of that model, including our own alternative argument. Section 3 summarizes the data used in the paper. Section 4 presents the paper's main empirical results. In Section 5 we discuss the main alternative interpretation of our results – that prejudice is correlated with unmeasured racial skill differences – and present the results of our attempts to deal with this problem by directly controlling for skill quality differences and through the use of instrumental variables. Section 6 concludes.

2. Theoretical Overview

Predictions from Becker's (1957) seminal work on employer prejudice are the foundation of the empirical analysis we conduct in this paper. In this section, we briefly review the key results from his basic model. We then discuss criticisms and extensions of the basic model that have appeared in the literature, including a brief discussion of our own theoretical argument which points out a peculiar, implicit assumption that has appeared in most previous analyses and which shows that discriminatory wage gaps in Becker's framework can survive the effects of long-run competition.

Becker's Basic Employer Discrimination Model

Throughout his analysis, Becker assumes a perfectly competitive environment, with production that is constant returns to scale. To focus attention on demand side considerations, white (a) and black (b) workers are assumed to be perfect substitutes in production. Employers are assumed to be white, and are differentially racially prejudiced. Becker represents prejudice as a distaste, or aversion, for cross-racial contact. Thus, an employer's utility, V_i , depends both on his profit and on the number of blacks he employs, with each black worker he hires bringing him disutility of $d_i \geq 0$. Employer utility is therefore

$$V_i = \pi_i - d_i L_b, \quad (1)$$

where $\pi_i = f(L_a + L_b) - w_a L_a - w_b L_b$ is the employer's profit, w_a and w_b denote white and black wages, respectively; and $f(\cdot)$ is the constant returns to scale production function. Employers choose white and black labor (L_a and L_b) to maximize (1).

The utility maximizing choices, L_a^* and L_b^* satisfy the conditions:

$$\begin{aligned} f'(L_b^* + L_a^*) - w_a &\leq 0, & \text{with equality if } L_a^* > 0 \\ f'(L_b^* + L_a^*) - w_b - d_i &\leq 0, & \text{with equality if } L_b^* > 0 \end{aligned} \quad (2)$$

Condition (2) says that if an employer hires a particular type of labor at all, he does so up to the point where its marginal product equals its marginal cost, as assessed by the employer. For white labor this marginal cost is the wage w_a ; for black labor it is the wage of blacks, plus the employer's prejudice, or $w_b + d_i$. This result captures the essence of Becker's insight: prejudice causes an employer to behave *as if* black workers' monetary wages are higher than they actually are. Since the two types of labor are perfect substitutes, (2) implies that an employer hires only white workers if his prejudice is such that $w_a < w_b + d_i$, and hires only black workers otherwise. These hiring rules mean that market pressures tend to sort prejudiced persons away from the object of their prejudice: for any given level of wages, blacks are hired by the least prejudiced employers in the market, while whites are hired by the most prejudiced.

Equilibrium in the short-run requires that the markets for white and black workers clear, at equilibrium wages w_a^* and w_b^* . Assuming a sufficiently smooth distribution of prejudice, some employer will be just indifferent between hiring black and white workers. The prejudice of this "marginal discriminator", d_i^* , is, in fact, equal to the equilibrium racial wage gap, since his being indifferent about the race of workers he hires at equilibrium wages implies

$$w_a^* = w_b^* + d_i^* . \quad (3)$$

Employers more prejudiced than the marginal discriminator hire only whites; those less prejudiced than d_i^* hire only blacks and the markets for both black and white workers clear in equilibrium.

Becker's simple, competitive framework yields a number of sharp yet subtle predictions about prejudice and equilibrium racial wage gaps. The first is that the equilibrium racial wage gap is *not* determined by the average level of prejudice among all employers. The foregoing shows that black wages are determined by the prejudice of the most prejudiced employer with whom blacks interact in equilibrium. Since blacks are, in utility terms, cheapest to the least prejudiced employers, they are hired first by these employers. The most prejudiced employer who hires blacks – the marginal discriminator – has, in general, less prejudice than the mean prejudice among all employers in the market. Indeed, as seen in Figure 2, the equilibrium wage gap might be 0 if there are many prejudiced employers in the market, so long as the supply of blacks and the

distribution of prejudice among employers are such that blacks can all be hired by totally unprejudiced employers.

The foregoing suggests that if racial wage gaps were empirically related to the average and the marginal level of prejudice among employers, Becker's model predicts that only for the latter measure should there be a systematic relationship. Two factors frustrate efforts to conduct this empirical test. One is that even if the complete distribution of prejudice among employers were known, it is impossible to know *ex ante* which employer is the marginal. Becker's original discussion *does* suggest one simple measure for the prejudice of the marginal employer that should hold under particular conditions. Specifically, if firms are of equal size, and if p is the fraction of blacks in the workforce, the marginal employer's prejudice will be the p^{th} percentile of the employer prejudice distribution. Of course, this measure of the marginal is an approximation and only applies under very specific conditions. Furthermore, this measure necessarily conflates the two things that determine the marginal – the number of blacks in the workforce and the distribution of prejudice among employers. Empirically, it might be attractive to separately assess the impact of each of these components of the marginal on the equilibrium wage difference.

Becker's model yields sharp predictions about how wage gaps should be separately related to these two variables. The effect of an increase in the number of blacks on equilibrium wages, all else equal, is clear. Holding constant the distribution of employer prejudice, an increase in the number of black workers means that blacks will, in equilibrium, have to be sorted to ever more prejudiced employers, so that the marginal employer is ever more prejudiced. This is illustrated in Figure 2 by a shift of the relative supply curve from S_1 to S_2 , reducing relative wages from 1 to R . The equilibrium wage gap should thus be increasing in the number of blacks in the market, holding constant employer prejudice.

The effect of a *ceteris paribus* increase in employer prejudice is more subtle, and depends on where in the distribution of prejudice that increase comes from. Since market pressures sort blacks to the least prejudiced employers first, holding constant the number of blacks and increasing the prejudice of the *most* prejudiced employers in the market should not affect who the marginal employer is, and thus should not affect the equilibrium wage gap. On the other hand, an increase in the prejudice of the *least* prejudiced employers while holding the number of black

workers constant may make the marginal employer more prejudiced, and thus increase the equilibrium wage gap. Such an increase in prejudice among those likely to be the marginal employer is illustrated in Figure 2 by the rotation of the relative demand curve from ABD to ABD' , thereby reducing relative wages from R to R' . An increase in prejudice among the most prejudiced would cause the relative demand curve to rotate beginning beyond the point where it intersects supply. In general, higher levels of prejudice in the right tail (higher percentiles) of the employer prejudice distribution should not affect equilibrium wages, whereas because of the sorting of blacks towards less prejudiced employers predicted by the model, relative black wages should fall as prejudice increases in the lower percentiles of the employer prejudice distribution.

Jointly, these predictions – (a) that the marginal matters more than the average prejudice for relative wage differences, (b) that the number (or fraction) of blacks in the workforce is negatively related to racial wage gaps, holding prejudice constant, (c) that prejudice in the right tail of the employer prejudice distribution should *not* matter for racial differences, while higher prejudice in the left tail of the prejudice distribution should affect racial wage gaps, and (d) that the mechanism which generates these patterns is the tendency of the market to segregate blacks from the most prejudiced whites – are the key results forthcoming from Becker’s basic employer prejudice discrimination model. Strangely, to our knowledge there have been no previous empirical tests of these predictions in the literature. We test for each of these predictions, with some modifications to be made clear below, in the empirical work to follow.

Extensions, Criticisms of Standard Becker Prejudice Model

Before turning to our empirical analysis in the next section, it is useful to briefly discuss some directly relevant criticisms and extensions of Becker’s model that have appeared in the literature. Beginning with Becker himself but most famously articulated by Arrow (1972, 1973), many authors have suggested that the equilibrium and predictions discussed above cannot hold in the long run under conditions of perfect competition. Central to their argument is the fact that in the short-run equilibrium discussed above, the firms less prejudiced than the marginal discriminator earn higher profits than their more prejudiced counterparts who choose to employ more expensive white labor. Since capital can move freely in the long run under perfect competition, unprejudiced firms will expand at the expense of their more prejudiced counterparts. Competition thus seems to ensure that prejudiced firms are driven out of the market in the long run. Racial wage gaps

deriving from employer prejudice disappear in the long run, and the predictions about an empirical relationship between wage differences and prejudice discussed above do not hold.

The notion that employer prejudice is “driven out of the market” in the long run remains a staple of most textbook treatments of the employer prejudice model.⁵ Yet, several authors have shown that prejudicial tastes can lead to persistent racial wage gaps in models with some form of either imperfect information or imperfect competition. In his discussion of the basic Becker model, Arrow argued that if there were adjustment costs or some other “non-convexity”, relatively unprejudiced employers need not expand at the expense of their more prejudiced counterparts. Recent work showing that prejudice can survive in the long run when there are adjustment costs include Lang, Manove and Dickens (2000) and Black (1995). Black shows that wage gaps resulting from prejudice can persist if there is costly search rather than the full information of the competitive model.⁶ In another departure from the original Becker model, Goldberg (1982) models racial sentiment not as a *distaste* for blacks but instead as nepotism, or favoritism towards whites. He shows that racial wage gaps, attributable to that type of prejudice, can survive in the long-run.⁷

The mechanisms that are the focus of Becker’s work – the tendency of market pressures to sort blacks away from the most prejudiced persons, and the fact that the ease with which the market can do this varies inversely with the number of blacks in the market – are also found in these extensions to his model. The extensions differ from Becker’s original formulation in that they typically assume either some sort of imperfect competition or imperfect information or else characterize racial prejudice slightly differently. We argue here that the original Becker predictions can be shown to survive in the long run even in a competitive environment, with racial preferences represented essentially the same way as Becker’s original. We summarize the argument verbally here, and offer a short formal sketch in the Appendix.

⁵ That employer prejudice cannot hold in the long run under perfect competition assumptions is found in most labor economics textbooks, and review pieces on the subject. See Ehrenberg and Smith (2003), Altonji and Blank (1999), and Lang (2007).

⁶ In Black’s model, the inability of black workers to costlessly locate prejudiced employers causes their prejudiced employers to enjoy monopsony power over them, with the result that black wages are relatively lower in equilibrium. Importantly, in order for prejudiced employers with higher costs to remain in business in the long-run despite their lower profits, Black assumes that entrepreneurial skill is scarce so that discriminators can pay these costs out of rents.

⁷ Under Goldberg’s formulation, nepotists are willing to pay, out of the return on their capital or the return on their own labor, for the utility they receive from employing whites. The only difference between nepotism and some other form of consumption is that nepotism cannot be purchased in a separate market.

The essence of the conventional criticism of Becker's original model is that prejudiced employers shut down or are driven out of the market because they sacrifice profit by remaining open and hiring more expensive whites. But a prejudiced employer in the long run has two options. On the one hand, he can remain in business and express his racial prejudices by paying more in wage costs to hire an all white workforce. As the standard criticism says, this choice involves a loss in money profit. On the other hand, the prejudiced employer can shut down, and not sacrifice money profit. If he were to shut down, however, the employer would become a *worker* at another (possibly unprejudiced) firm, and would have to interact with that firm's black workers as a fellow employee. As a prejudiced individual, he must therefore take account of the racial composition of the firm at which he would work when considering whether to shut down his own firm. To argue that a prejudiced employer necessarily shuts down because of the higher monetary costs of hiring white workers is to say that he ceases to be prejudiced once he assumes the role of employee. In other words, the conventional criticism of Becker implicitly makes what we regard as the unrealistic and peculiar assumption that racial prejudice is not *portable* across the different roles that a person might play in the market.

Under the more realistic assumption that agents take their prejudices against interacting with blacks across the different roles they play in the labor market – that the prejudiced employer becomes a prejudiced employee after shutting down his firm – it does *not* necessarily follow that prejudiced employers shut down in the long run. Prejudiced employers must consider both the monetary returns and the racial composition of their outside option. Thus the equilibrium racial composition of firms and the ability of the market to segregate workers are key factors in determining whether prejudiced individuals remain as employers. Indeed, if agents are assumed to have the same level of racial prejudice irrespective of the labor market role that they play, every person who is a prejudiced employer in the short run will shut down in the long run only so long as the market is sufficiently segregated by race that each can be assured of finding employment as a worker in a firm in which no blacks are employed. That is, precisely as argued by Becker in his original model, the degree to which employer prejudice is related to racial wage gaps in equilibrium varies inversely with the degree to which the market can be segregated by race.

There are two interesting consequences of assuming that racial prejudice is portable across labor market roles. The first is that doing so combines two types of prejudice models – employer and employee prejudice – which Becker formally analyzed in distinct models. The second related

point is that since the labor market roles of “employer” and “worker” are endogenously chosen by market participants, to the extent that racial wage gaps are a function of racial prejudice among whites, what matters is not the distribution of prejudice among people who happen to be employers at a point in time, but rather the distribution of prejudice among *all* whites.

3 Data Summary

In this section we describe the data on racial prejudice used in the empirical analysis to test the predictions of the prejudice model. We summarize racial prejudice across different regions in the U.S., across different population subgroups, and the evolution of racial prejudice over the past thirty years

Overview of Prejudice Data

The General Social Survey (GSS) is our source of data on racial prejudice. We use data from multiple waves (1972-2004) of the GSS. In many survey years, this nationally representative data set elicited responses from survey questions about matters that are clearly strongly related to racially prejudiced sentiments. “Prejudice” is a nebulous construct, and the various questions posed in the survey over the years touch on the different dimensions along which racist sentiments might be manifest. In fact, over the approximately 30 years of GSS data used in the paper, respondents answered some twenty-six different questions relating to some aspect of racial feeling. A different subset of the full questions was asked each year, with no particular question asked in each year of the survey, and with much variation in the total number of times a given question appears.

One concern with the GSS questions is that whereas all of the questions touch on something having to do with racial feeling, some of them probably also reflect some other type of preference. The types of questions for which this concern is strongest are questions having to do with government policy. For example, in many years respondents were asked whether they believed that “the government was obligated to help blacks”. Persons with no feelings of racial animus whatever might still respond in the negative to this question because of their views about the appropriate role of government. Given the paper’s aim, it is imperative that we focus on responses that have to do exclusively with racial sentiment. We therefore do not use any of the five questions touching on government policy and race in our analysis. For reasons discussed

below, in part of the analysis we need to focus on a consistent set of prejudice questions from one year to the next.⁸ When necessary, we therefore focus on four questions that *jointly* appear most frequently in the survey.

Appendix Table 1 lists the GSS variable abbreviation and a summary for each of the full set of 26 racial prejudice questions asked in the survey. Most questions are in bold face. Those not in bold are the government policy questions we always exclude, and the shaded questions are the set of four race questions asked consistently across the survey. Throughout, we use responses from whites aged 18 and older, and recode responses so that higher values correspond to more prejudiced answers.⁹

Much of our analysis involves comparing levels of prejudice across individuals and across geographic areas. To render these comparisons feasible it is obviously necessary that we somehow combine the disparate prejudice responses into a uni-dimensional prejudice index. We do this in two simple steps. First, we normalize the various prejudice responses using the 1977 report. Formally, let d_{it}^k denote respondent i 's response in year t to the particular prejudice question k . For each dimension of prejudice k and for each individual i we create a normalized individual response in year t given by

$$\tilde{d}_{it}^k = \frac{d_{it}^k - E[d_{i,77}^k]}{\sqrt{\text{Var}(d_{i,t_k^*}^k)}} \quad (4)$$

where t_k^* is the first year in which the prejudice question k was asked in the GSS. The normalized measure thus subtracts off from individual responses to each question the mean of the response to that question in 1977, and divides by the standard deviation of answers measured in the first year the question was asked.^{10 11} These normalized responses, which are all measured on

⁸ Although they are not included in our analysis, we find that the government policy questions exhibit the same time series properties and variation across states and regions as the prejudice questions we use.

⁹ In most cases, this recoding is straightforward (e.g. those who would not vote for an otherwise qualified black person for president are more prejudiced than those who would). In some cases the ordering of responses is less clear (e.g. those who think the federal government is spending too much improving the conditions of blacks may not be prejudiced; they may think the federal government is spending too much on everything). However, in each case we think it is clear which response was *meant* to denote greater prejudice.

¹⁰ We normalize by the standard deviation in the first year the question was asked rather than, say, the overall standard deviation, because we want to avoid a mechanical relationship between trends in responses and the weight the question receives in the overall aggregate.

the same scale, are then aggregated into a one dimensional aggregate prejudice index for individual i in year t by taking their average in the year, or

$$D_{it} = \sum_k \tilde{d}_{it}^k / K_t \quad (5)$$

where K_t is the number of prejudice questions asked in year t .

To test the predictions of the prejudice model across different geographic regions, we use several measures of the prejudice among whites in a community. We compute aggregate measures of prejudice from $\tilde{D}_{it}$, which is the residual from a regression of D_{it} on a full set of year dummies. The first measure of aggregate prejudice in a community – denoted “average” prejudice – is simply the mean across all years of $\tilde{D}_{it}$ for whites in a particular geographic area. Another set of measures captures prejudice at different percentile points in the overall prejudice distribution in a state. We use the 10th, 50th and 90th percentiles of the distribution of $\tilde{D}_{it}$ within a state. Finally, we create an approximation to the prejudice of the “marginal” white discriminator in the distribution. As noted in the theory overview, under specific conditions, the marginal is well approximated by the p^{th} percentile of the distribution of prejudice where p is the fraction of the labor force that is black.^{12,13} Notice, since we use year-residualized individual prejudice $\tilde{D}_{it}$ to compute all of the community-level prejudice indices, the downward secular trend in individual-level prejudice that we document below does not lead us to disproportionately weight recent

¹¹ Two prejudice questions (HELPBLK and RACOPEN in the table) were not asked in 1977, but were asked in both prior and subsequent years. A linearly interpolated mean is subtracted for these variables instead of the 1977 mean.

¹²For the percentile measures (including the marginal), we use an aggregate index $\tilde{D}_{it}$ which is constructed from the 4 consistent questions asked in the GSS. We do not use the full 26 questions when creating these percentile measures because the aggregate measure in that case is an average over responses to different numbers of questions in different years. The variance of the aggregate index based on the full 26 questions will tend mechanically to be higher in years in which a smaller number of questions happened to be asked. Statistics based on the tails of the distribution would disproportionately measure prejudice in those years in which a relatively small number of questions were asked.

¹³All of the results presented below are virtually unchanged if the average is computed using only the four questions that are consistently asked together throughout the GSS.

observations when calculating indices that measure prejudice in the left tail of the distribution (i.e. the 10th percentile or the marginal).¹⁴

Summarizing Patterns and Trends in Prejudice in the GSS

Before turning to the regressions that are our main focus, we present some initial results about prejudice to provide the reader a sense of the data. Table 1 presents a series of regressions showing how key demographic traits are related to an individual's level of prejudice in the GSS. The table shows results for four outcomes: the individual aggregate prejudice index, and un-normalized average responses to three of the specific prejudice questions from which the uni-dimensional aggregate index is constructed. The regressions are estimated on a pooled GSS sample across all years, and each regression controls for fixed year and state fixed effects. The standard errors presented in the table are clustered at the level of the state.

The regression for the aggregate prejudice index shows that prejudice exhibits a sharp age gradient, in that older whites are significantly more likely to report prejudiced sentiments.¹⁵ To give a sense of the magnitude of the estimate, a decade-older white in the sample is on average about 13 percent of an individual-level standard deviation (which is about 0.7) more prejudiced. The regression also shows that higher-educated persons and females are significantly less prejudiced. Four years of education are associated with a reduction in prejudice of about a third of a standard deviation, and females are around 5 percent of a standard deviation less prejudiced than males. When all of variables are jointly controlled for, women and the more highly educated are both found to be significantly less prejudiced, and the strong age gradient remains significant. These results for the specific prejudice measures are generally very similar to those for the aggregate index, with the notable exception that white women appear slightly more prejudiced than white men with respect to their opposition to inter-racial marriage.

Table 2 summarizes patterns of prejudice across regions in the U.S. The first three columns of the table show average responses in the various Census divisions to the same three race sentiment

¹⁴ The results are virtually unaffected if aggregate prejudice measures are computed based on D_{it} rather than $\tilde{D}_{it}$. Some of the point estimates are larger in specifications that use D_{it} as the base distribution, but a one-standard-deviation change in each case implies virtually the same effect.

¹⁵ The regressions control for year effects, but the usual problem restrains us from making conclusions about whether the age gradient is driven by true age effects or cohort effects.

questions depicted in Table 1.¹⁶ As described above, higher values indicate greater racial prejudice. The table shows that the specific prejudice measures are very highly correlated.¹⁷ By each of the measures, racial prejudice is most severe in the southeastern portion of the country and least severe in New England and in the West. Prejudice is greatest in the East South Central division (AL, KY, MS, TN), and next greatest in the South Atlantic (DE, D.C., FL, GA, MD, NC, SC, VA, WV), and West South Central (AR, LA, OK, TX) divisions. Prejudice is least severe in New England (CT, ME, MA, NH, RI, VT), and in the Pacific (AK, CA, HI, OR, WA) and Mountain (AZ, CO, ID, MT, NV, NM, UT, WY) divisions.

The third and fourth columns of the table show the means for the aggregate prejudice index – first over the entire sample of whites in a region, and then over a group of “high skilled” (college-graduate) whites, who might be more likely to employers than the average person. Not surprisingly, the table reveals the same pattern of prejudice across regions as do the separate prejudice questions. This is true for both the sample of all whites and the sample for the highly skilled, although the means show that, in any division, the highly skilled are less racially prejudiced than their less educated counterparts.¹⁸ The difference in average prejudice across the various divisions is substantial. For example, between the East South Central and New England Census divisions the difference is on the order of 0.8 of an *individual-level* standard deviation. To put this difference in perspective, the median East South Central respondent has the same aggregate prejudice as the 81st percentile respondent from New England. The median-prejudiced New England respondent would be at the 26th percentile of the East South Central prejudice distribution. A graphical description of the cumulative distribution of prejudice across the various divisions is presented in Appendix Fig 1. The fifth column of the table shows the fraction of the labor force that is black in each Census division. On the whole, the numbers indicate that blacks live disproportionately in regions of the country where racial prejudice is most severe, on average. This pattern may indicate that prejudice is caused in part by cross-racial contact and by competition for economic resources.

¹⁶ The possible answers to the questions in the first two columns are 1 (yes), or 0 (no), whereas there are four possible answers to the question about residential segregation in the third column: 1 (disagree strongly), 2 (disagree slightly), 3 (agree slightly), and 4 (agree strongly).

¹⁷ To conserve space we obviously cannot show the same numbers for all of the prejudice questions. Reassuringly, the basic patterns shown in Tables 1 and 2 are found for other measures as well.

¹⁸ The values for the univariate average prejudice measure are generally negative. This is because this measure is based on normalized measures described above, and all measures of prejudice have been declining since 1977.

Using the individual prejudice measures, we earlier showed graphically that black wages are relatively lower where prejudice is higher. The regression results presented in the last two rows repeat this finding, and also show that racial wage gaps are larger where whites are more prejudiced according to the aggregate prejudice index.¹⁹ At best, these patterns are merely suggestive about the causal connection between prejudice and wages suggested by Becker's model. Also suggestive is the fact that, although the correlation is weaker, the fraction of a labor market that is black is also correlated with the level of the black-white wage gap. The analysis below tests more carefully whether the patterns of prejudice and wage gaps fit the subtle patterns suggested by the theory. In that work, we address such issues as the fact that prejudice indices might be correlated with unobserved regional differences in productivity between black and white workers.

Before turning to the main analysis, we show how prejudice has varied over the time period we study. Figure 3 shows trends in responses to the nine most commonly asked GSS prejudice questions (excluding, as noted earlier, those having to do with government policy), averaged across the entire sample of whites. As described above, each question is normalized so that the mean response in 1977 is zero, and the standard deviation in the first year it was asked is 1. The average response among whites to each question has declined steadily over the past 30 years. The figure reveals a general downward trend in each reported measure of racial prejudice, although there is substantial variation in the magnitudes of these declines. Declines are very small for an objection to sending one's children to school with blacks, and quite for large responses about whether blacks should "push where they are not wanted".

Figure 4 shows trends over time in the uni-dimensional aggregate prejudice measure, plotted separately for each Census division. Two key things should be noted about the figure. First, the decline in measured prejudice has been widespread. Between 1977 and 1996, measured prejudice declined in each of the nine Census divisions. Second, the relative ranking of average levels of prejudice across regions has been constant over time. New England, and the Pacific and Mountain divisions were the least prejudiced regions throughout the period under study, while the East South Central and South Atlantic divisions were consistently the most prejudiced, according to our index. Third, there appears to have been some convergence across Census divisions in the

¹⁹ Because our aggregation method treats all prejudice responses equally, the measures used in the paper do not make ad hoc judgments about which specific prejudice questions better reflect true underlying racial prejudice. In effect, our use of equal weights takes a "hands-off" approach to the available data.

level of measured prejudice between the early 1970's and the mid-1990's, though the amount of convergence is sensitive to the choice of endpoint. On the whole, these time series patterns suggest that the data only permit us to speak confidently about differences in prejudice across regions, rather than differences over time in the relative declines in prejudice across spatial areas.

4 Base Empirical Results

To estimate the relationship between relative black wages and region-level measures of prejudice we merge the prejudice indices described above with CPS data. We combine the May monthly supplement from 1977 and 1978 with the Merged Outgoing Rotation Group (MORG) files from 1979 to 2002.^{20,21} The sample includes full-time black and white males aged 16 to 64.²² Our basic specification could be estimated in two alternative ways. One option would be to estimate an OLS regression of log wages on education, a quadratic in potential experience, race-specific year dummies, a black dummy, the average prejudice for the state, and the interaction between the particular prejudice measure and a black indicator variable. The coefficients of interest would be the estimated effect on the interaction terms. In the case of the average, for example, an estimated negative coefficient would indicate how much lower black wages are relative to whites in states with higher average prejudice. The drawback of this method is that because the various prejudice indices only vary at the level of the state, this procedure might underestimate standard errors, even if the regressions were clustered at the level of the state.

An alternative method, which deals with this potential standard error concern, is executed in two steps. First, we estimate the residual black-white wage gap in each state. Specifically, we estimate by OLS the log wage regression described above, but leave out the prejudice index and include a separate black dummy variable for each state. The estimated effects on each of these black dummy variables become the dependent variable in the second step, which is weighted by the precision with which we estimate the state wage gap in the first step. In this second step regression, one or more of the labor market prejudice indices are the independent variables of

²⁰ See e.g. Lemieux (2006) or Autor, Katz and Kearney (2005) for a discussion of the merits of the May and MORG files for measuring wages. We follow Autor et. al.'s sample restrictions, dropping those with real hourly wages below the real value of the 1982 minimum wage or with nominal wages above top code levels. Top-coded responses are replaced with 1.5 times the top-code value. We thank David Autor for sharing his programs with us.

²¹ The analysis below requires state-level wage gaps. State is not consistently reported in the CPS until 1977. For the earlier analysis at the Census division level, we also use data from the 1973-1976 May CPS.

²² Results are similar using a sample of all black and white males aged 16 to 64 with positive earnings.

interest. The second step regression has approximately 45 observations and therefore produces conservative standard error estimates.²³ Reassuringly, the results are substantively almost identical to those from the one-step procedure described above.

Table 3 presents the main results. The regressions in the table assess whether the adjusted racial wage gaps across states vary with alternative measures of prejudice and with the racial makeup of those states in a manner predicted by the Becker prejudice model reviewed earlier. The table and most that follow report point estimates and standard errors for the average, median, 10th, and 90th percentiles of the white prejudice distribution; the fraction of the labor market that is black; and the approximation to the “marginal level of prejudice described earlier, and given by the p^{th} percentile of the prejudice distribution where p is the fraction of the state labor force that is black.

Column 1 shows the pairwise relationship between the black-white wage gap and the average prejudice among whites in a state. These state-level results show a weaker relationship between wage gaps and average prejudice than the earlier results shown at the Census division level. The point estimate is negative, suggesting that Black wages are relatively lower in states with higher average prejudice, but the relationship is not statistically significant. In contrast, the pairwise relationship between the black-white wage gap and the prejudice of an approximation of the “marginal” white in the labor market, shown in column 2, is striking. As the Becker model predicts, the regression shows that states in which the marginal white is more prejudiced have substantially larger racial wage gaps. The standard deviation of “marginal” prejudice across states is 0.139. The estimated effect therefore implies that a one-standard deviation increase in prejudice is associated with relative black wages that are 0.028 log points lower – about a 23 percent effect relative to the mean residual wage gap across states. The result is estimated precisely enough to marginally reject a zero correlation at the 5-percent significance level.

Becker’s model suggests more than that these variables should independently be systematically related to wage gaps. If prejudice matters as the model predicts, then measured wage gaps should be related to the marginal and not the average when the two measures are included jointly in a regression. Column 3 of the table reports the results of this specification. The results show that, holding the average level of prejudice constant, the estimated effect of marginal prejudice among

²³ States are dropped because they are not separately identified in the GSS. These states are small and tend to have very few blacks.

whites is negative, strongly statistically significant, and larger in absolute value than the pairwise estimate in column 2. That the correlation of the wage gap with the marginal is stronger than with the average is precisely as Becker's prejudice model predicts. The sorting of blacks away from the most prejudiced white employers implies that the wage gap is determined on the margin by the racial tastes of individuals in the lower portion of the prejudice distribution.

Variation in the marginal level of prejudice among whites across states comes from two sources: differences in the number of blacks across the states, and differences in the distribution of prejudice across states. As discussed earlier, the prejudice model implies that the fraction black should have an *independent* negative effect on relative black wages. If the model is correct, holding the level of prejudice constant, the larger the number of blacks in a region the more likely it is that blacks are sorted to ever more prejudiced employers. Thus, whereas the results in column 3 show that wage gaps load onto the average and not the marginal, a better assessment of how the model's predictions fit the data is forthcoming from the regression in column 4 in which we relate racial wage gaps simultaneously to the average, the marginal, *and* the fraction black in the state.

The results show that the fraction of blacks in the state is strongly negatively related to the wage gap, as predicted. At the same time, we find that when the number of blacks is controlled for, the point estimate for the marginal is about 60 percent as large, but remains very strongly statistically significant and of the sign predicted by theory. That the estimated effect for the marginal is smaller than the estimates in columns 2 and 3 is as expected since, as noted in the discussion of figure 2, variation in the marginal is driven partly by the number of blacks in the state. The fact that we continue to find a strongly statistically significant negative effect of the marginal after directly controlling for percent black shows suggests that variation in prejudice in the part of the distribution near the p^{th} percentile of the state's prejudice distribution is negatively related to wages. Finally, the regression shows that the estimate of the effect of average prejudice is positive, and not statistically significant.

One criticism of our interpretation – that the estimate of the marginal conditional on fraction black reflects the effect of variation in prejudicial tastes – is that the measure of the marginal might pick up direct non-linear effects of fraction black. If so, controlling for fraction black linearly will not remove all the variation of the marginal that is attributable to state racial makeup. To assess the importance of this issue, we estimate the effect of the marginal controlling for

fraction black non-linearly. These various estimates are shown in Appendix Table 3. Column 1 presents the base specification, which includes average and marginal prejudice, along with linear fraction black. Columns 2-4 include increasingly higher-order polynomial terms in fraction black. The marginal is significant in each case, about 9 percent smaller in magnitude and fairly consistent as we allow fraction black to go from quadratic to quartic. Column 5 shows results from a less parametric specification in which we control for a set of dummy variables for different categories of fraction black. Again, the estimated effect of the marginal is statistically significant and is almost identical in magnitude to the base specification with a linear fraction black control. These results suggest that the marginal is not picking up non-linear direct effects of fraction black. Rather, it appears that the variation in the marginal after controlling for percent black represents variation in prejudice.

Taken together, the results in the first four columns of Table 3 are strongly consistent with the predictions of Becker's prejudice model. However, these results do not constitute the sharpest possible associational tests of the model's prediction. For one thing, our measure of the marginal is an approximation to what the true marginal likely is in a state. Moreover, as discussed in the theory overview, Becker's model yields sharp predictions about how wage gaps should be affected by prejudice in a particular part of the prejudice distribution but not by prejudice at other points.

In column 5 we relate state level wage gaps to the level of prejudice at the 10th, 50th and 90th percentiles of the overall prejudice distribution. The results are striking. We find a large, significant, negative effect of the prejudice at the 10th percentile. The estimated effects imply that the median and 90th percentiles of the prejudice distribution have no effect on relative black wages. The results in the last column of the table are even more striking. When we control for the fraction of the state that is black in addition to the three percentile points in the prejudice distribution, the effect of the 10th percentile prejudice is still negative and strongly statistically significant. Prejudice at the median and 90th percentiles is again estimated to have no effect on state level racial wage gaps. Finally, the fraction of the state that is black is in this regression strongly negative and statistically significant.²⁴

²⁴ As discussed above, we argue that it is the distribution of prejudice among whites that determines relative black wages and not prejudice among persons who, at a point in time, happen to be playing the role of labor market employer. Nonetheless, we estimate a version of Table 5 but with prejudice measured among the set of people who are likely to be employers – those with high skill. When the various measures are

These results are strongly consistent with the predictions of the prejudice model. It is worth emphasizing that they say much more than that “black wages are lower in places where whites are more prejudiced”. The fact that in a regression in which they are jointly controlled for, prejudice among the least prejudiced whites in an area matters for relative black wages while that of the most prejudiced whites has no effect at all, combined with the fact that the fraction of the state that is black varies in a negative and statistically significant way with relative black wages suggests that only a mechanism of the sort first outlined by Becker can reconcile these facts.²⁵ Indeed, the higher variance found in prejudice at the right tail and illustrated in Appendix Figure 1 might have led one to speculate that right tail prejudice would be found to matter more empirically. That it does not, makes the results all the more striking.²⁶

In the next section, we address the two concerns that might lead us to question the conclusion that the patterns in Table 3 reflect what could be termed the causal effect of prejudice and racial composition on racial wage gaps. Before turning to that analysis, we briefly address one potential question about the interpretation.

This question has to do with the fact that regional differences in answers to the GSS questions may not reflect differences in racial prejudice at all. Instead, it could be the case that racial feelings are the same everywhere, but there are regional differences in people’s willingness to be candid about their underlying racial sentiments. In our view, this argument begs the question of why people in different places feel differentially obliged to hide their underlying racial sentiment. One could regard this willingness to be candid about racially insensitive feelings as *itself* indicative of the thing we have called prejudice. More importantly, notice that if there were no

constructed from the sample of whites with at least some college training, we find results very similar to those shown in Table 3. The results are presented in Appendix Table 2.

²⁵ In the wage context, only prejudice in the left tail matters because racial interactions are mediated by a price mechanism operating in a market context. We would not expect a similar prediction in situations where the racial outcomes are determined not by a market but rather by some other mechanism, such as a vote. In fact, as one example, we find that the log public welfare spending per black in a state (measured in 1978, 1987 or 1996) is significantly negatively related to the 90th percentile of prejudice in the state but not to the median or 10th percentile. To the extent that policies that benefit blacks relative to whites are determined by the median level of prejudice or prejudice in the right tail, it is more likely that the patterns we see in the wage regressions are driven by variation in prejudice rather than unobserved omitted factors.

²⁶ We have experimented with including other percentile points of the prejudice distribution, such as the 25th and 75th percentiles. Although not shown, these regressions show the same thing as the results in column 5 of Table 4: only prejudice in the left tail of the prejudice distribution (such as the 10th or 25th percentile) varies in a negative and statistically significant way with relative black wages. Prejudice in the right tail consistently has no effect on the racial wage gap.

content to people's claims about their racial feelings, the GSS racial prejudice measures would basically be noise and there should be no systematic relationship between stated prejudice and wage differences, as our various analyses find.

5. Extensions: Accounting for Endogenous Preferences and Racial Composition

On the whole, the results presented in Table 3 are strongly supportive of all of the predictions of the prejudice model of wage discrimination. But can a causal interpretation be given to these results? The obvious specific concern here is that regions with more severe measured prejudice or higher numbers of blacks also have other unobserved characteristics that negatively affect black wages more than white wages. Though the regressions we have discussed thus far control for a rich set of fixed effects, and the most obvious suspects, such as differences in education levels, there are always other possibilities. One example is that some determinant of labor market outcomes, such as unobserved school quality (not quantity) could be correlated with prejudice levels.²⁷ In fact, this reduced school quality could be a direct *result* of prejudice, as the important historical example of segregated schools illustrates. A second possibility is that unobserved racial skill differences, as might be measured by the black-white test score gap, could be correlated with prejudice. The patterns we have documented could result from unobserved skill differences of this sort even if there were no direct effect of prejudice on wages in the labor market. The other type of potential endogeneity concern in Table 3 is the fact that the fraction of a population that is black may be related to some unobserved determinant of wages. Potentially magnifying both of these concerns is that in the regressions we have presented thus far, most of the variation in the data comes from comparisons between the Southern states and the rest of the country. Anything peculiar about labor market institutions or skill levels in the South relative to the rest of the country could thus possibly explain our results.

Of course the ideal way to address any potential concern about the endogeneity of prejudice and of the fraction of a state that is black would be to find valid instruments for these variables. In the case of the various percentile measures of the prejudice distribution, we do not have and cannot conceive of an instrument which could, for example, generate plausibly exogenous variation in the 10th percentile of the prejudice distribution but not the 90th, or vice versa. To deal with any potential endogeneity of these measures, we therefore adopt the approach of collecting

²⁷ Notice that it would not cause a bias if school quality were relatively lower for *both* blacks and whites since we include state effects, unless the effect of school quality on wages differed by race.

information on the factors for which there is the greatest concern that prejudice is correlated at the state level – measures of skill quality in the state.

Tables 4 and 5 present the results from directly controlling for racial skill differences. In Table 4 we re-estimate the base results in Table 3, but add controls for the state-level racial difference in mean National Assessment of Educational Progress–Long Term Trend (NAEP-LTT) math and reading scores.²⁸ These measures of racial skill differences have the expected sign, and the wage gap appears to be more significantly related to gaps in reading than math skills. Somewhat surprisingly though, we find that adding these controls for the full set of states in the sample leave the main qualitative results virtually unchanged. We continue to find that the measure of the marginal level of prejudice is more closely related to the racial wage gaps than the average, that the lower parts of the prejudice distribution matter while the right tail matters little if at all, and that the fraction of the state that is black is always strongly negatively related to relative black wages.

In Table 5, we use another indicator of latent skill differences by race. In their paper on relative school quality, Card and Krueger (1992) collect information on the student/teacher ratios for various cohorts of students in 18 southern states. We use these reported ratios for the cohort born between 1940 and 1949 as a measure of possible racial skill differences.²⁹ Importantly, the analyses reported in this table are restricted to Southern states. There are two panels in the table. In the upper panel, we simply re-estimate the set of regressions shown previously, but only on the sample of Southern states. Notice that the basic patterns found in the full sample are found within this specific set of states as well. Wage gaps are more closely related to marginal than average prejudice; the marginal is significantly negatively related to relative black wages, even conditional on average prejudice and fraction black; and wage gaps are significantly related to the

²⁸ The NAEP-LTT is a standardized test administered to U.S. students and is designed to measure trends and cross-sectional patterns in educational performance. It is sometimes referred to as “The Nation’s Report Card”. Since 1971, students of age 9, 13, and 17 have been randomly selected and tested in mathematics and reading. Each subject is tested approximately every four years. The content tested by the NAEP-LTT has remained unchanged so that comparisons over time can be made. We use a restricted-use version of the NAEP-LTT that has state identifiers and individual data. For each student, we transform math and reading scores into a z-score (mean zero, standard deviation of one) and then compute state-level differences in average z-scores by race, computed over the full sample of years (1971-2004 for reading, 1978-2004 for mathematics). The NAEP-LTT was also administered in mathematics in 1973, but we do not have state identifiers for those data.

²⁹ Card and Krueger (1992) report black and white pupils per teacher for each of four cohorts. The results are substantively the same if the relative pupils per teacher for the 1910-1919, 1920-29 or 1930-39 are instead used as controls.

10th percentile of the prejudice distribution but not the median or the 90th, conditional on fraction black. The bottom panel repeats the analysis above, but adds the Card/Krueger measure of black-white relative school quality to each regression. The strong similarity between these results and those presented earlier is quite striking. The results indicate that our results do not appear to be driven by un-observed differences between the South and the rest of the country. They also show that when we control for a widely regarded measure of school quality (and thus skill) differences within the South, key results remain basically qualitatively unchanged.

We turn next to the possible endogeneity of the fraction black variable in the various regressions. Table 6 presents instrumental variables estimates in which we instrument for the fraction black during the time period we study with the fraction black in the state in 1920. For columns 2, 3, and 4, the results are for regressions in which we instrument the marginal using an alternative measure of the marginal based on the fraction black in the state in 1920 and the distribution of prejudice from the GSS sample period.³⁰ These historical population shares are obviously correlated with the number of blacks currently living in a state. In effect, the IV strategy isolates and uses only that variation in the current racial makeup of a state attributable to the historical accident of where blacks were located eight decades ago. To the extent that this historical racial makeup of states is independent of current determinants of wages, the estimates are purged of the endogeneity concerns from recent migration.

The table shows that this measure for the marginal yields results that are strikingly similar to the baseline results. We stress that we cannot reject the possibility that a state's historical racial makeup is correlated in *some* way with factors that determine wages nearly a century later but the results raise confidence that the results we are estimating really reflect the effect of the marginal discriminator being drawn from a higher point in the distribution of prejudice rather than some correlation between unmeasured productivity and current racial composition. The results in column 6 are probably even more convincing. In these regressions we include the different percentile measures, and simply instrument for the fraction black. Again, the results are qualitatively almost identical to the main results presented earlier.

³⁰ The first-stage coefficients on the marginal are 0.339 and 0.370 with t-statistics of 16.31 to 10.11 for the specifications in columns 2 and 3, respectively. The first-stage coefficient on fraction black is 0.582 with a t-statistic of 7.99 for the specification in column 6.

In summary, we believe that these results taken together strongly support the results presented in the previous section, and suggest that the relationship we document between relative wages and both the prejudice in a state and the fraction of the state that is black derives causally from the mechanisms described in Becker's prejudice model.

Examining the Role of Workplace Segregation

Before concluding, we present one final set of results assessing the relationship between workforce racial segregation and racial wage gaps. In the review of the model, we stressed that a key point of Becker's reasoning is that the degree to which racial wage gaps are observed in equilibrium will depend on the extent to which the market can be segregated by race. The only data that we could find about the extent to which whites interacted with blacks at work – the dimension of segregation most relevant for thinking about the effect of contact discrimination – is from the 2000 wave of the NELS-88. NELS respondents were asked what fraction of their co-workers was of the same race as them.³¹ The NELS population is relatively young in 2000, but this information about workplace segregation will be closely related to what is true of whites overall if there is no dramatic difference in the distribution of employment for whites of different ages. The mean response to this question, by state, for whites measures what the segregation literature calls the white exposure to black index. We use this as our measure of workplace racial integration.

The first column of Table 7 shows that, consistent with the theoretical discussion, the racial wage gap is significantly larger the more integrated whites are with blacks at work. In the second and third columns, we add, in turn, controls for the average and marginal prejudice in the state. In the specification that controls for average prejudice, racial integration at work remains strongly statistically significant. Very strikingly, column 3 shows that the estimated effect of the segregation measure falls by about half and is statistically insignificant when only the marginal is included in the regression. That there is no such comparable reduction when only average prejudice is added to the regression suggests that, as we have argued throughout, who the marginal discriminator turns out to be is determined to a large degree by how much segregation is possible in the labor market. In columns 5 and 6, we include the various percentiles of the prejudice distribution and find that black relative wages are significantly negatively related to the

³¹ Specifically, people were asked: "What percentage of the people in your present/most recent workplace are of the same race as you?"

degree of workplace integration conditional on the full distribution of prejudicial tastes and the fraction black workforce in the state. The fact that the effect of workplace integration falls more when the marginal is controlled for than when the percentiles of the prejudice distribution and percent black are included argues further that a mechanism linking wages and prejudice, operating through the market's ability to separate races at work, determines the racial wage gap, as we have argued throughout.

6. Conclusion

Our primary goal in this paper has been to empirically test the sharp but subtle predictions about the relationship between community-level prejudice and racial wage gaps implied in Becker's seminal model of employer discrimination. To our knowledge, no previous paper in the large literature on discrimination has attempted to do this.

We motivate our analysis by first reviewing the short-run version of Becker's classic model, from which the predictions that are central to our analysis come. We then discuss criticisms by Arrow and others, arguing that those predictions cannot hold in the long run under standard competitive assumptions. We point out that this interpretation has become the standard discussion in most textbook discussions of employer prejudice, despite the fact that various authors, emphasizing alternative non-competitive assumptions, have shown that prejudice can indeed survive in the labor market in the long run. We briefly outline a simple theoretical framework that reaches the same conclusion, but does so by keeping within Becker's original framework of a competitive market and with preferences of the sort originally discussed by Becker.

Using rich data on prejudice from multiple years of data from the General Social Survey, we summarize both the cross sectional variation and trends over time in racial prejudice among whites. We document significant variation in prejudice across different regions of the country. We also show that while reported prejudice has declined significantly everywhere over the past thirty years, the magnitude of that decline has been fairly consistent across regions, though there is some evidence of convergence.

In our main empirical analysis, we test for, and confirm, a series of key predictions from the standard Becker model. First, we show that racial wage gaps are much more closely related to the level of prejudice of the "marginal" person in the distribution, than they are to average levels of

prejudice. Even more striking, we show that it is only prejudice in the left tail of the prejudice distribution that seems to matter for wage gaps; wages do not vary at all with the prejudice of the most prejudiced persons in a state. This is precisely as the Becker model predicts. Finally, we show that the fraction of a labor market that is black is strongly negatively related to racial wage gaps – again, just as the prejudice model predicts. We conduct a variety of robustness tests, including instrumental variables estimates and regressions in which we control directly for skill differences across regions that might be an alternative explanation for our results. Our key results are robust to all of these extensions.

In our view, the paper’s various results point to an important role for racial prejudice in wage determination for minorities. Clearly, much more work, both on the theoretical front and with respect to empirical analysis, needs to be done in order for us to have a better sense of the ways in which prejudice operates and the effect it has on wages. For example, we have analyzed a particularly simple form of racial animus: an aversion to cross-racial contact. In this we follow Becker, who assumes that this is the form that racial prejudice takes. However, racial animus can take other forms that might be relevant for wage determination. Explicit theoretical analysis of alternative formulations of prejudice is an obvious next step for future work. In addition, an important area for future work would be to assess how racial prejudice, perhaps through its effect on wages, affects other important outcomes like migration or occupational choice. It is also worth stressing that the results we have documented about racial prejudice explain only a portion of racial wage differences and therefore do not imply that other mechanisms posited in the literature such as statistical discrimination and human capital differences do not also matter importantly for racial wage determination.

Finally, whereas we believe that the empirical evidence we have presented is strongly suggestive of an important role for racial prejudice on wage, in a manner consistent with theory, we have been careful to stress that absent quasi-experimental evidence, care must be taken in attaching causal interpretations to our estimates. Future work, in which scholars find suitable instruments for individual or community prejudice is an obvious next step on the empirical front.

Bibliography

Aigner, Dennis J. and Glenn G. Cain (1977). "Statistical Theories of Discrimination in Labor Markets." Industrial and Labor Relations Review, 30, pp. 175-187.

Altonji, Joseph G. and Rebecca M. Blank (1999). "Race and Gender in the Labor Market," in Orley C. Ashenfelter and David Card, eds. Handbook of Labor Economics, Volume 3C. New York: Elsevier, pp. 3143-3259.

Altonji, Joseph G. and Charles R. Pierret (2001). "Employer Learning and Statistical Discrimination." Quarterly Journal of Economics, 464 (February), pp. 313-350.

Arrow, Kenneth (1972). "Some Mathematical Models of Race in the Labor Market", in A. H. Pascal ed. Racial Discrimination in Economic Life. Lexington, MA: Lexington Books, pp. 187-204.

Arrow, Kenneth (1973). "The Theory of Discrimination," in O. A. Ashenfelter and A. Rees, eds. Discrimination in Labor Markets. Princeton, N.J.: Princeton University Press, pp. 3-33.

Autor, David H., Lawrence F. Katz and Melissa S. Kearney (2005). "Trends in U.S. Wage Inequality: Re-assessing the Revisionists." NBER Working Paper No. 11627.

Becker, Gary S. (1957, 1971). The Economics of Discrimination. Chicago: The University of Chicago Press, 2nd edition.

Bertrand, Marianne and Sendhil Mullainathan (2004). "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." American Economic Review 94(4), September, pp 991-1013.

Black, Dan A. (1995). "Discrimination in an Equilibrium Search Model." Journal of Labor Economics 13(2), pp. 309-334.

Borjas, G.J. and S.G. Bronars (1989). "Consumer Discrimination and Self-employment." Journal of Political Economy 97(3) (June), pp. 581-606.

Cain, Glen G. (1986). "The Economic Analysis of Labor Market Discrimination: A Survey," in O. Ashenfelter and R. Layard, eds. Handbook of Labor Economics Volume One. Amsterdam: North-Holland, 693-785.

Card, David, Alexandre Mas and Jesse Rothstein (2006). "Tipping and the Dynamics of Segregation in Neighborhood Schools." mimeo. Princeton University.

Chay, Kenneth Y. (1995). "The Impact of Federal Civil Rights Policy on Black Economic Progress: Evidence from the Equal Opportunity Act of 1972." Princeton University Industrial Relations Section Working Paper 346.

Cutler, David, Edward Glaeser and Jacob Vigdor (1999). "The Rise and Decline of the American Ghetto." Journal of Political Economy 107(3) (June), pp. 455-506.

Doeringer, P.B and Michael J. Piore (1971). Internal Labor Market and Manpower Analysis. Lexington: Heath.

Donohue, John J. III and James Heckman (1991). "Continuous versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks." Journal of Economic Literature 29 (December), pp. 1603-43.

Fryer, Roland and Steve Levitt (2004). "The Causes and Consequences of Distinctively Black Names." Quarterly Journal of Economics 119(3), pp 767-805.

Goldberg, Matthew S. May (1982). "Discrimination, Nepotism, and the Long-Run Wage Differential." Quarterly Journal of Economics 97(2), pp 307-319.

Granovetter, Mark. (1995). Getting a Job: A Study of Contacts and Careers. 2nd ed. Chicago: University of Chicago Press.

Heckman, James J. (1998). "Detecting Discrimination." Journal of Economic Perspectives 12(2), (Spring), pp. 101-116.

Lang, Kevin, Michael Manove and William Dickens. (2005). "Racial Discrimination in Markets with Announced Wages." American Economic Review, 95, 1327-40.

Nardinelli, Clark and Curtis Simon (1990). "Customer Racial Discrimination in the Market for Memorabilia: The Case of Baseball." The Quarterly Journal of Economics, 105 (3) (Aug), pp. 575-595.

Neal, Derek A. and William R. Johnson (1996). "The Role of Premarket Factors in Black-White Wage Differences." Journal of Political Economy 104 (5), pp. 869-895.

O'Neill, June (1990). "The Role of Human Capital in Earnings Differences between Black and White Men." Journal of Economic Perspectives 4 (Fall), pp. 25-45.

Oettinger, Gerald (1996). "Statistical Discrimination and the Early Career Evolution of the Black-White Wage Gap." Journal of Labor Economics 14(1), pp. 52-78.

Phelps, Edmund S. (1972). "The Statistical Theory of Racism and Sexism." American Economic Review, 62, pp. 659-61.

Rosen, Sherwin (1986). "The Theory of Equalizing Differences." In O. Ashenfelter and R. Layard, eds., Handbook of Labor Economics.

Smith, James P. and Finis Welch (1989). "Black Economic Progress After Myrdal." Journal of Economic Literature, 27 (June), pp. 519-64.

Smith, James P. (1993). "Affirmative Action and the Racial Wage Gap." American Economic Review, 83(2) (May), pp. 79-84

Theory Appendix

The key implication of the argument made by Becker, Arrow and others is that in the long run the return to capital must be the same across all firms, regardless of the firm's prejudicial tastes or racial composition. Any capital owner who does not earn this competitive return can sell his capital to an unprejudiced person, to the benefit of both parties.³² Suppose therefore that all employers rent capital at a price P_k . Suppose further all persons who could be employers have the option of selling their labor as workers, receiving the wage w_a . If there is a racial wage gap an unprejudiced employer receives a utility payoff from being an employer of

$$\Pi(w_a^*, w_b^* | d_i = 0) = f(L_b^*) - w_b^* L_b^* - P_k. \quad (6)$$

Of course, since in equilibrium any unprejudiced person could replace this unprejudiced employer if the payoff from being an employer were higher than the wage that person could earn as a worker, it must be the case that:

$$w_a = f(L_b^*) - w_b^* L_b^* - P_k^* \quad (7)$$

where P_k^* is the long run equilibrium price of capital that must be paid by any person who is an employer.

In the short run, employers more prejudiced than the marginal hire only white workers, while those less prejudiced than the marginal hire less expensive blacks. Let the difference in short-run money profits between the two types of employers be³³

$$\Delta = f(L_b^*) - w_b^* L_b^* - [f(L_a^*) - w_a^* L_a^*] > 0 \quad (8)$$

In the long run, when he has to pay the competitive price of capital, an employer who is more prejudiced than the short run marginal discriminator receives a monetary payoff from being an employer given by

$$f(L_a^*) - w_a^* L_a^* - P_k^* = w_a - \Delta \quad (9)$$

Since this prejudiced employer could alternatively sell his labor on the competitive market as a worker, he would experience a net monetary loss of Δ by operating as an employer in the long run.

Although this argument is not usually presented formally, it is clearly the essence of the conventional argument, going back to Becker and Arrow. However, economic agents make

³² One reading of Becker (1957) and Arrow (1972) is that they take employers to be the owners of the firm's capital. One reason prejudiced capital owners leave the market is that their outside option does not involve market interaction with blacks. In contrast, we take the term "employer" to combine two roles: a capital owner or renter and a supervisor who makes hiring, firing, and wage decisions and who is himself a labor input.

³³ This difference in short-run profit (the essence of conventional criticisms of the standard Becker model) must exist by the following reasoning. Since in the short-run a totally unprejudiced employer is indifferent about the race of the workforce he hires, and since such an employer *could* have hired the workforce L_a^* at the equilibrium wage w_a^* , but chose to hire the L_b^* at the wage w_b^* instead, it *must* be true that firms more prejudiced than the marginal discriminator earn less profit in the short run than unprejudiced employers.

decisions on the basis of utility rather than money profits. Thus, in deciding whether to stop being an employer in the long run, a prejudiced employer with an all white workforce compares (9) not to his alternative wage, but rather to the *utility* he would get as a worker. He shuts down in the long run if

$$w_a - \Delta \leq w_a - \kappa_i \tilde{l}_b, \quad (10)$$

where $\tilde{l}_b$ is the number of blacks with whom he would work as a fellow worker, and κ_i is the disutility he gets from cross racial contact with fellow employees who are black. The left hand side of (10) is the utility a prejudiced employer receives by continuing to operate as an employer in the long run. It is the residual dollar profits from his firm after he has paid the competitive rate for capital. Since he chooses not to hire any black workers, he suffers no disutility from interacting with them as an employer. The right hand side is the same person's utility if he were to become a worker: it equals his wage *minus* his disutility from interacting with any black co-workers were he to shut down and become an employee at another firm.

The conventional argument that prejudice is driven from the market in the long-run under competition implicitly assumes that $\kappa_i = 0$ – that is, that a person who so dislikes interacting with blacks as an employer that he avoids hiring them altogether, somehow has no aversion against interacting with them as a fellow employee. We argue that it is much more reasonable to assume $\text{Cov}(d_i, \kappa_i) > 0$, so that racial prejudice is *portable* across the roles of employer and co-worker. The simplest form of portability is when $d_i = \kappa_i$. If this holds then a prejudiced person will definitely shut down and become a worker in the long run only if the profit he forgoes by staffing his workforce with more expensive white workers is greater than his disutility from interacting with black *as a co-worker* at any new job he might take, or

$$\Delta \geq d_i \tilde{l}_b. \quad (11)$$

With portable racial prejudice (11) holds for certain in the long run only if the market can be segregated enough by race so that all prejudiced workers are able to work in firms with no blacks. Any impediment to segregation in the real world, including mandated racial quotas in firms, costs of searching, or imperfect substitutability of high and low skilled workers in production, means that the employer shutting down his firm may expect to encounter some fellow black workers at his new job ($\tilde{l}_b > 0$). The conventional notion that employers engaging in discriminatory hiring are driven out of the market by competition in the long run need not hold.³⁴

³⁴ The point can perhaps be seen most clearly if we consider the possibility that because of constant returns to scale, a single alternative employer, presumably with no prejudice, expanded enough to account for all production but that done by a single prejudiced employer. In that case, the prejudiced employer's alternative option were he to shut down would be to *definitely* work with black co-workers. If his prejudice is portable, he would derive disutility from this option and should be willing to incur some dollar profit loss to prevent it.

Table 1: Demographic Traits and Individual Level Prejudice

	Aggregate Index of Individual Prejudice				Not Vote for a Black for President			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age/10	0.090 (0.003)			0.071 (0.002)	0.030 (0.003)			0.022 (0.002)
Education		-0.057 (0.003)		-0.047 (0.002)		-0.021 (0.002)		-0.018 (0.002)
Female			-0.038 (0.006)	-0.068 (0.005)			-0.016 (0.006)	-0.026 (0.005)
No. Obs.	35,757	35,780	35,864	35,684	16,441	16,463	16,491	16,416
R-squared	0.16	0.17	0.11	0.20	0.07	0.08	0.05	0.09
	Support a Law Against Interracial Marriage				Whites Have Right to Segregate Neighborhoods			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Age/10	0.070 (0.003)			0.054 (0.002)	0.137 (0.005)			0.105 (0.005)
Education		-0.044 (0.003)		-0.035 (0.002)		-0.086 (0.004)		-0.070 (0.004)
Female			0.030 (0.006)	0.009 (0.005)			0.017 (0.013)	-0.030 (0.011)
No. Obs.	23,368	23,378	23,433	23,319	15,294	15,304	15,337	15,264
R-squared	0.20	0.21	0.12	0.26	0.17	0.17	0.11	0.20

Note: Table reports coefficient estimates from individual level OLS regressions of measures of prejudice on demographic traits. Four regressions are reported for each dependent variable, denoted above the results. In addition to the regressors listed in the table, regressions control for state and year effects. Standard errors are corrected for clustering at the state level.

Table 2: Prejudice and Wages across Census Divisions

	Support law against interracial marriage	Would not vote for Black for President	Whites have right to segregate neighborhood	Average Prejudice Index		B-W Wage Gap	
				Overall	High Skilled	% Black	
E. Sou. Central	0.504	0.330	2.356	0.167	-0.123	14.2	-0.281
South Atlantic	0.377	0.236	2.187	0.007	-0.341	16.9	-0.249
W. Sou. Central	0.306	0.210	2.011	-0.090	-0.358	9.7	-0.212
E. Nor. Central	0.245	0.146	2.007	-0.168	-0.484	6.9	-0.143
W. Nor. Central	0.243	0.152	1.930	-0.201	-0.515	2.2	-0.117
Middle Atlantic	0.203	0.133	1.919	-0.210	-0.465	8.6	-0.202
Mountain	0.159	0.104	1.642	-0.359	-0.560	1.7	-0.149
New England	0.149	0.085	1.647	-0.375	-0.604	2.4	-0.172
Pacific	0.132	0.098	1.628	-0.378	-0.547	4.5	-0.109
Dep. Var.: Unconditional B-W Wage Gap, Bivariate OLS Coefficients (std. err.):							
	-0.404 (0.102)	-0.617 (0.158)	-0.181 (0.056)	-0.263 (0.068)	-0.338 (0.078)		
Observations	9	9	9	9	9		
R-squared	0.69	0.69	0.60	0.68	0.73		

Note: The top panel reports sample means for each of the nine Census divisions. The possible answers to the questions in the first two columns are 1 (yes), or 0 (no), whereas there are four possible answers to the question about residential segregation in the third column: 1 (disagree strongly), 2 (disagree slightly), 3 (agree slightly), and 4 (agree strongly). The bottom panel reports coefficients and standard errors for bivariate OLS regressions of the black-white wage gap in the Census division on the measure of prejudice reported in the respective column.

Table 3: Estimated Relationship Between Racial Prejudice of Whites in a Labor Market and Black-White Relative Wages

<i>Measure of Prejudice among All Whites</i>	Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d.) of Black-White Wage Gap Across States: -0.123 (0.044))					
	(1)	(2)	(3)	(4)	(5)	(6)
Average	-0.036 (0.030)		0.097 (0.029)	0.050 (0.033)		
Marginal		-0.213 (0.040)	-0.328 (0.050)	-0.202 (0.068)		
10th Percentile					-0.212 (0.180)	-0.292 (0.125)
Median					-0.006 (0.062)	0.007 (0.043)
90th percentile					0.016 (0.029)	0.016 (0.020)
Fraction Black				-0.157 (0.062)		-0.304 (0.045)
State	45	45	45	45	45	45
R-squared	0.03	0.40	0.52	0.59	0.05	0.56

Note: Table reports coefficients (standard errors) from OLS regressions of residual state-level black-white wage gaps on various measures of prejudice among all whites. Residual black-white wage gaps are estimated using 1977-2002 May/ORG CPS data and control for education, a quadratic in experience, race-specific year effects and state effects. Data from 1973-1976 are dropped because CPS reports states in groups in those years. States are dropped if they are not sampled in the GSS in the years necessary to measure the marginal index of prejudice. The "marginal" is the p-th percentile of the prejudice distribution of the relevant population of whites, where p is the fraction of the population that is black. See text for details.

Table 4: Controlling for Test-score Differences by State

Dependent Variable: Residual Black-White Wage Gap in Market						
<i>(Mean (s.d.) of Black-White Wage Gap Across States: -0.123 (0.044))</i>						
<i>Measure of Prejudice among</i>						
<i>All Whites</i>	(1)	(2)	(3)	(4)	(5)	(6)
Average	-0.070 (0.028)		0.066 (0.037)	0.026 (0.039)		
Marginal		-0.196 (0.038)	-0.283 (0.061)	-0.171 (0.076)		
10th Percentile					-0.352 (0.190)	-0.336 (0.143)
Median					0.017 (0.057)	0.002 (0.043)
90th Percentile					-0.002 (0.030)	0.012 (0.023)
Fraction Black				-0.155 (0.069)		-0.265 (0.052)
White-black difference in NAEP Math	-0.024 (0.040)	-0.034 (0.033)	-0.034 (0.032)	-0.042 (0.030)	-0.030 (0.042)	-0.055 (0.032)
White-black difference in NAEP Reading	-0.141 (0.039)	-0.084 (0.031)	-0.047 (0.037)	-0.026 (0.036)	-0.149 (0.041)	-0.047 (0.037)
States	39	39	39	39	39	39
R-squared	0.34	0.55	0.59	0.65	0.37	0.65

Note: Table reports coefficients (standard errors) from specifications similar to those in table 3, adding controls for racial difference in standardized test scores. The test-score control is the difference in average z-score from the NAEP Long Term Trend Math (1978-2004) and Reading (1971-2004).

Table 5: Restricting to Southern States and Controlling for School Quality

A:						
<i>Measure of Prejudice among All Whites</i>	Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d.) of Black-White Wage Gap Across States: -0.143 (0.038))					
	(1)	(2)	(3)	(4)	(5)	(6)
Average	0.008 (0.040)		0.107 (0.024)	0.013 (0.034)		
Marginal		-0.206 (0.053)	-0.319 (0.043)	-0.165 (0.057)		
10th Percentile					-0.383 (0.246)	-0.300 (0.116)
Median					0.012 (0.065)	0.019 (0.031)
90th Percentile					0.062 (0.037)	-0.005 (0.020)
Fraction Black				-0.194 (0.058)		-0.325 (0.046)
States	18	18	18	18	18	18
R-squared	0.00	0.49	0.78	0.88	0.23	0.84
B:						
Average	0.043 (0.038)		0.108 (0.024)	0.017 (0.035)		
Marginal		-0.187 (0.068)	-0.295 (0.051)	-0.158 (0.060)		
10th Percentile					-0.025 (0.260)	-0.287 (0.155)
Median					-0.055 (0.063)	0.017 (0.038)
90th Percentile					0.067 (0.032)	-0.004 (0.022)
Fraction Black				-0.188 (0.061)		-0.321 (0.060)
White-black pupil- teacher ratio	0.266 (0.105)	0.048 (0.104)	0.063 (0.069)	0.032 (0.055)	0.281 (0.117)	0.011 (0.083)
States	18	18	18	18	18	18
R-squared	0.30	0.49	0.79	0.88	0.47	0.84

Note: Panel A of the table restricts the sample to the 17 Southern states plus Missouri, the set of states for which Card and Krueger (1992) collected school quality measures. Panel B also restricts to the same 18 states, and adds a control for the white-to-black pupil teacher ratio in the state as reported in Card and Krueger (1992).

Table 6: Two-stage Least Squares Estimates Using Fraction Black in 1920 as an Instrument

<i>Measure of Prejudice among All Whites</i>	Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d.) of Black-White Wage Gap Across States: -0.123 (0.044))					
	(1)	(2)	(3)	(4)	(5)	(6)
Average	-0.036 (0.030)		0.127 (0.033)	0.070 (0.044)		
Marginal		-0.204 (0.044)	-0.401 (0.063)	-0.251 (0.102)		
10th Percentile					-0.212 (0.180)	-0.309 (0.128)
Median					-0.006 (0.062)	0.010 (0.044)
90th Percentile					0.016 (0.029)	0.016 (0.021)
Fraction Black				-0.175 (0.094)		-0.367 (0.055)
State	45	45	45	45	45	45
R-squared	0.03	0.40	0.50	0.57	0.05	0.54

Note: Table reports coefficients (standard errors) from 2SLS regressions of residual state-level black-white wage gaps on various measures of prejudice among all whites. Fraction black in the state estimated in the 1920 census is used as an instrument for the contemporaneous fraction black, and an alternative marginal based on the fraction black in 1920 is used as an instrument for the contemporaneous marginal.

Table 7: Estimated Effect of Prejudice of Whites and Workplace Segregation on Black-White Relative Wages

Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d) of Black White Wage Gap Across States: -0.123(0.044))							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Non-White Co-Workers of White Workers in State	-0.224 (0.093)	-0.226 (0.092)	-0.113 (0.079)	-0.041 (0.077)	-0.102 (0.074)	-0.252 (0.101)	-0.172 (0.071)
Average		-0.037 (0.028)		0.092 (0.031)	0.029 (0.036)		
Marginal			-0.196 (0.041)	-0.315 (0.055)	-0.150 (0.077)		
10th Percentile						-0.083 (0.177)	-0.199 (0.123)
Median						-0.066 (0.063)	-0.035 (0.044)
90th Percentile						0.032 (0.028)	0.027 (0.020)
Fraction Black					-0.183 (0.064)		-0.287 (0.043)
Observations	45	45	45	45	45	45	45
R-squared	0.12	0.15	0.43	0.53	0.61	0.18	0.62

Note: Measure of Workplace Integration is estimated from white respondents' answer to a question in NELS-88 about racial makeup of co-workers. See text for details.





Figure 1:
Relationship between Census Division Black-White Wage Gap
and Two Prejudice Related Questions from the GSS




Figure 2:
The Relationship between Racial Tastes and the
Relative Wages and Relative Supply of Blacks and Whites

Note: The figure shows how the equilibrium ratio of black to white wages responds to three sets of market conditions. When the relative supply of black workers is small relative to the number of unprejudiced employers, as is the case when supply is as depicted by S_1 , the marginal discriminator is unprejudiced and there is no racial wage gap in equilibrium. Holding constant the distribution of racial preferences among employers, a shift out in the relative supply of black workers (from S_1 to S_2) requires that more prejudiced employers hire blacks, and the ratio of black to white wages falls from 1 to R . Holding constant the relative supply of black workers, an increase in prejudice among employers likely to be the marginal discriminator (which causes the relative demand curve to rotate from ABD to ABD'), further reduces the equilibrium ratio of black to white wages to R' .




Figure 3:
Trends in responses to GSS prejudice questions

Note: Full descriptions of questions are listed in appendix table 1.





Figure 4: Trends in prejudice by Census division

Note: Figure plots average over time by Census division using the four prejudice questions jointly asked most frequently in the GSS.



Appendix Figure 1:

Cumulative distribution of the individual prejudice index, by Census division.

Appendix Table 1: GSS questions used to measure prejudice

AFFRMACT	Do you oppose a preference in hiring and promotion?
BUSING	In general do you favor the busing of black and white children from one school district to another?
CLOSEBLK	In general, how close do you feel to blacks?
FEELBLKS	In general, how warm or cool do you feel towards blacks?
HELPBLK	Agree? The government is obligated to help blacks.
NATRACE	Agree? We are spending too much money improving the condition of blacks.
RACAVOID	If you were driving through neighborhoods in a city, would you go out of your way to avoid going through a black section?
RACCHNG	If you and your friends belonged to a social club that would not let blacks join, would you try to change the rules?
RACDIN	How strongly would you object if a family member brought a black friend home for dinner?
RACJOB	Do you think blacks should have as good a chance as anyone to get any kind of job, or do you think white people should have the first chance at any kind of job?
RACMAR	Do you think there should be laws against marriages between blacks and whites?
RACMAREL	How would it make you feel if a close relative of yours were planning to marry a black?
RACMARPR	Agree? You can expect special problems with marriages between blacks and whites.
RACOBJCT	If a black with the same income and education as you have, moved into your block, would it make any difference to you?
RACOPEN	Would you vote for a law that says a homeowner can refuse to sell to blacks, or one that says homeowners cannot refuse to sell based on skin color?
RACPEERS	Aggregation of three questions about whether you would object to sending your kids to a school that had few/half/most black students.
RACPRES	If your party nominated a black for President, would you vote for him if he were qualified for the job?
RACPUSH	Agree? Blacks shouldn't push themselves where they're not wanted.
RACQUIT	If yes to RACCHNG: If you could not get the rules changed, do you think you would resign from the club, even if your friends didn't?
RACSCHOL	Do you think white students and black students should go to the same schools or separate schools?
RACSEG	Agree? White people have the right to keep black people out of their neighborhoods and blacks should respect that right.
RACSUBGV	Do you think the city government in white suburbs should encourage black people to buy homes in the suburbs, discourage them, or leave it to private efforts?
RACSUBS	Do you oppose voluntary (religious/private business) efforts to integrate white suburbs?
RACSUPS	Agree? You can expect special problems with black supervisors getting along with workers that are mostly white.
RACTEACH	Agree? A school board should not hire a person to teach if that person belongs to an organization that opposes school integration.
WRKWAYUP	Agree? Italians, Jews and other minorities overcame prejudice and worked their way up. Blacks should do the same without special favors.

Note: Table lists each of the 26 questions from the GSS used to measure prejudice. The four questions shaded in gray were asked in the 1972, 1977, 1985, 1988, 1989, 1990, 1991, 1993, 1994 and 1996 waves of the GSS. We use these four questions to construct the prejudice indices that vary within region over time, as well as the indices of the marginal discriminator's prejudice and the various percentile measures of prejudice. In all but one case, the variable name is the same as the one listed in the GSS codebook. RACPEERS is based on three variables (RACFEW, RACHAF, RACMOST), which ask "Would you yourself have any objection to sending your children to a school where [a few/half/most] of the children are blacks?" Some of the descriptions are the verbatim questions asked in the survey, while others are paraphrased to save space. Questions were asked in various years of the GSS.

Appendix Table 2: Estimated Relationship Between Racial Prejudice of High-Skilled Whites in a Labor Market and Black-White Relative Wages

<i>Measure of Prejudice among High-skilled Whites</i>	Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d.) of Black-White Wage Gap Across States: -0.123 (0.044))					
	(1)	(2)	(3)	(4)	(5)	(6)
Average	-0.079 (0.036)		-0.005 (0.034)	-0.031 (0.031)		
Marginal		-0.293 (0.057)	-0.288 (0.066)	-0.106 (0.074)		
10th Percentile					-0.095 (0.127)	-0.193 (0.089)
Median					-0.064 (0.064)	0.010 (0.045)
90th Percentile					0.009 (0.024)	-0.021 (0.017)
pb100				-0.226 (0.059)		-0.309 (0.046)
States	45	45	45	45	45	45
R-squared	0.10	0.38	0.38	0.55	0.08	0.57

Note: Table reports coefficients (standard errors) from specifications similar to those in table 3, but where the prejudice indices are computed only using high-skilled (i.e. those with at least a college education) whites.

Appendix Table 3: Confirming the Effect of the Marginal is Robust to Allowing Fraction Black to Enter Non-Linearly

<i>Measure of Prejudice among All Whites</i>	Dependent Variable: Residual Black-White Wage Gap in Market (Mean (s.d.) of Black-White Wage Gap Across States: -0.123 (0.044))				
	(1)	(2)	(3)	(4)	(5)
Average	0.080 (0.029)	0.054 (0.034)	0.054 (0.034)	0.054 (0.035)	0.050 (0.033)
Marginal	-0.201 (0.075)	-0.185 (0.076)	-0.183 (0.077)	-0.183 (0.079)	-0.202 (0.068)
Fraction Black	X	X	X	X	
Fraction Black ²		X	X	X	
Fraction Black ³			X	X	
Fraction Black ⁴				X	
Fraction Black in Categorical Dummies					X
States	45	45	45	45	45
R-squared	0.58	0.59	0.60	0.60	0.59

Note: Table reports OLS regressions that test whether the index of marginal prejudice is robust to non-linear controls for fraction black. Column 1 of the table reports the base specification that is also shown in column 4 of table 3. Columns 2-4 add polynomials in fraction black of increasingly higher order. Column 5 replaces the linear fraction black control with a set of mutually exclusive dummies based on categories of fraction black.