

Prejudice and the Economics of Discrimination

Kerwin Kofi Charles
University of Chicago and NBER

Jonathan Guryan
University of Chicago GSB and NBER

February 14, 2007

Abstract

This paper re-examines the role of prejudice in market discrimination against racial minorities. We begin by reviewing Becker's (1957) model of employer discrimination, outlining a number of direct implications we test empirically in the second half of the paper. We then show that the widely accepted theoretical result from Becker (1957) that competition ultimately forces prejudiced employers to shut down can be reversed with two reasonable assumptions. First, we explicitly characterize an employer as being both the residual-claimant owner of the firm's capital, as in Becker (1957), and as a manager who makes hiring, firing and wage-setting decisions and who is himself a labor input. Second, we argue that a prejudiced employer's tastes would be taken with him to any other role he might occupy in the labor market – a prejudiced employer who shuts down and takes a new job would be a prejudiced worker at that job. We consequently show that whether prejudiced employers shut down in the long-run depends on the expected racial composition of their co-workers at other firms, through its effect on their outside options. Discrimination can therefore survive if something prevents firms from segregating perfectly by race – a result we illustrate with the example of imperfect substitution in production. We argue that Becker's *employee* discrimination model should be viewed as a general model of discrimination, with the roles of employers and employees determined endogenously.

The second part of the paper presents empirical evidence on racial prejudice and its relationship to black-white wage gaps across regions of the U.S, something not previously done in the large discrimination literature. We discuss trends and cross-sectional patterns in racial prejudice using data from the General Social Survey. We then show that as predicted by our characterization of Becker's model, the self-employed are more prejudiced than the average worker. Since most self-employed have no employees, we take this result to be evidence of market-induced segregation. Next, using data from the Current Population Survey from 1973 to 2002, we show that the black-white wage gap is greatest in the most prejudiced regions of the U.S. Finally, we show that, as predicted by Becker (1957), black-white wage gaps are more closely related to the prejudice of the "marginal discriminator" than to the average prejudice in an area.

The authors thank Gary Becker, John Bound, Charlie Brown and Kevin Murphy for helpful discussions regarding this work, seminar participants at Wharton, CUNY and Brown for their suggestions and Phil Wengerd for research assistance. Guryan thanks the University of Chicago GSB and the Industrial Relations Section at Princeton University for research support. All errors are our own.

1. Introduction

Becker's (1957) "The Economics of Discrimination" launched the formal analysis of labor market discrimination among economists, and provided a framework that has profoundly influenced all subsequent work on the subject. 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.

In the short-run version of the employer discrimination model, racial prejudice causes some employers to regard black workers as more expensive than they truly are. Despite being equally productive, the model shows that blacks might nonetheless receive lower wages in equilibrium. In the long run, however, as new, unprejudiced employers enter the market, the resulting competitive pressure forces erstwhile discriminating employers to shut down and leave the market. The notion that employer, taste-based discrimination cannot survive in the long run was further fleshed out by Arrow (1972), who memorably remarked that the employer discrimination model "predicts the absence of the phenomenon it was designed to explain."¹ Faced with this apparent limitation of taste-based models of employer discrimination, most recent theoretical work on discrimination has tended to ignore the role of prejudicial tastes among employers entirely as a source of racial wage gaps, emphasizing instead considerations like 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

But is the widely-held view that discrimination based on employer racial distaste cannot survive in the long run warranted? And, is the corresponding research focus on alternative explanations

¹ Arrow (1972), p. 192.

² Some recent empirical work offers tantalizing evidence that racial prejudice might matter for racial wage determination. For example, in an interesting and widely-known paper, Bertrand and Mullanaithan (2004) 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, possibly racial prejudice, might be at work.

to the virtual exclusion of analyses of racial prejudice justified? 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. Given this, plus the profound influence of Becker's conceptualization of prejudice, this paper re-examines the possible role of employer racial prejudice as a source of racial discrimination. We conclude that the commonly accepted idea that prejudice-based discrimination cannot survive in competitive markets is not theoretically true under a reasonable modification of the assumption about how prejudice likely operates. Further, we present substantial empirical evidence showing that racial prejudice, appears to matter empirically for racial wage gaps, contrary to the null under which most economists now operate.

Our analysis is broken into two parts. The first re-examines prejudicial tastes in a model of competitive equilibrium. This theoretical part of our analysis differs from Becker's original model and the analyses that followed in three simple but subtle ways. First, we clarify that the term employer refers both to the owner of the firm's capital—as Becker's analysis emphasizes—and to the manager residual-claimant who makes hiring, firing and wage setting decisions and who is a labor input himself. Next, we treat the terms "employer" and "co-worker" as referring less to distinct market actors than to particular *roles* played in the labor market. In other words, we endogenize the choice of whether to be an employer, or manager. Finally, we acknowledge explicitly that the prejudiced employer who shuts down his firm in the face of competition carries any racial animus he possesses into any other labor market role he plays: his prejudice is *role-independent*, or *portable*. That is, a prejudiced employer who shuts down becomes, in our framework, a *prejudiced* worker, concerned both about the wages he receives and the race of his new co-workers.

We show that the very reasonable assumption that an individual's racial prejudices are independent of the particular role he plays in the labor market dramatically changes the conclusions about when or whether prejudiced employers shut down their firms in the face of long-run competition. We show, in particular, that a prejudice-based model yields a long-run equilibrium in which discrimination is *not* necessarily driven out of the market. Because a prejudiced employer must consider the counterfactual racial composition of his co-workers when deciding whether to shut down his firm, the persistence of racial wage differences due to discrimination depends on the market's ability to segregate workers by race. Any hindrance to segregation may lead to long-run differences in wages attributable to prejudice. Put differently,

of the two market forces Becker emphasized for their role in reducing market discrimination—competition and segregation—we believe segregation should receive more attention than it has from economists. We formally show this result using the example of imperfect substitutability in production by skill level. Finally, since our analysis endogenizes an individual’s decision to be a worker or an employer, it makes clear that employer and employee discrimination which have hitherto been analyzed in two separate models can, and should, be treated in a single framework.

Having re-established a potential important theoretical role for racial prejudice in a standard competitive model of wage determination, the second part of our paper turns to an exercise never before conducted in the large literature on discrimination – formally assessing the empirical relationship between prejudice among whites and the racial wage gap. To measure prejudice, we use the rich information on racial sentiments available in multiple waves of data from the General Social Survey (GSS). We summarize racial prejudice among whites as a whole, and among white employers and high-skilled workers separately. After showing broad patterns of racial prejudice across regions, and summarizing how they have changed over time, we test one implication of the Becker’s theory that we feel is important: that more prejudiced workers should seek out firms with fewer black co-workers, rather than seek out a non-supervisory role. Consistent with this prediction, we show that the self-employed, who tend to have no employees or coworkers, are more prejudiced than the average worker, but that those in supervisory occupations are equally prejudiced to non-supervisors conditional on measures of skills.

We next formally assess the association between these prejudicial feelings and the wages of blacks and whites across different regions in the U.S. To conduct this part of our analysis, we merge data from the GSS with individual level data from the Current Population Survey (CPS). We perform a series of regressions relating individual wages to various measures of prejudice among whites in the census division. Generally, we find that the racial wage gap tends to be largest in census divisions where average reported prejudice among whites is highest. Though slightly less robust, we also find that black wages have decreased most in those divisions with the smallest reductions in racial animus over the past thirty years. These relationships are found whether average prejudice in the division is measured among whites overall, among white employers, or among high-skilled workers specifically. The standard prejudice discrimination model suggests that it is not the prejudice among all whites or among all white employers which determines the racial wage gap, but rather the prejudice of the employer who is indifferent between hiring a black and a white worker in equilibrium. We use a simple theoretically-based

argument, combined with information about the incidence of blacks in the division, to identify the (prejudice of) the marginal discriminator in the data. The results are quite striking and surprisingly robust, given the obvious measurement error inherent in identifying the marginal discriminator. Just as the model would predict, we show that the overall wage penalty experienced by blacks is largest where this marginal employer is more racially prejudiced. Despite the strong correlation between the various community prejudice measures studied, we show that in a “horse-race” of the various measures, most of the associations above load onto the marginal discriminator term rather than the average prejudice measures.

The absence of an instrument for prejudice in the regression analysis prevents us from arguing that the various associations we document are causal. However, the rich set of covariates for which we are able to control using the CPS data means that we are able to exclude some of the most obvious alternative explanations for the patterns we document – such as the notion that the most racist places in the country are also places where black are less skilled, at least along the dimension of years of completed schooling. And, we estimate broadly similar results, though slightly less robust, when we conduct analyses examining changes within a census division over time. Furthermore, the “horse-race” specifications show that black relative wages are lower in census divisions where the marginal discriminator is more prejudiced, even controlling for measures of average prejudice in the division. And using the measure of marginal discriminator implied by the model, black relative wages have increased most in the census divisions where the prejudice of the marginal discriminator has declined most, even controlling for changes in measures of average prejudice in the division.

On the whole, the results are broadly suggestive of an important role for racial prejudice among whites in explaining racial wage differences. The combination of these empirical results with the theoretical results from the first part of the paper showing that employer prejudice can theoretically lead to long run wage differences for blacks argues strongly for a reconsideration of the notion that racial prejudice has little to do with explaining observed wage and income differences by race.

The remainder of the paper proceeds as follows. The next section briefly reviews the short and long run equilibria in Becker’s employer racial distaste model. In Section 3 we present an analysis in which prejudices are treated as role independent, so that the racial prejudice of an employer is taken with him to any other role he occupies in the market – including especially that

of an employee at another firm. We present results assuming, in turn, perfect and imperfect substitutability of labor in production to illustrate how the latter assumption generates results different from the standard Becker analysis. In Section 4 we present empirical evidence of the distribution and evolution of racial prejudice among whites, the relative prejudice of supervisors and the self-employed, and the effect of community-wide racial prejudice on the relative wage received by blacks. Section 5 concludes.

2. A Review of Becker's (Dis)Taste Models of Discrimination

Basic Setup

In the simplest version of the employer distaste model, Becker assumes that black (B) and white (W) workers are perfect substitutes in production. Becker's key assumption is that an employer's utility depends on both the profits he receives and on the number of blacks he hires. Specifically, an employer i receives disutility d_i^{ER} from each black worker he employs. Let employer's utility, V_i , be given by

$$V_i = \pi_i - d_i^{ER} L_B. \quad (1)$$

In (1), $\pi_i = f(K, L_B + L_W) - rK - w_B L_B - w_W L_W$ is the employer's profit. The variables L_W and L_B refer to white and black labor; w_W and w_B denote black and white wages; r is the cost of capital; and $f(\cdot)$ is a production function that is constant returns to scale in capital and labor. We drop the i subscripts on the firm's choice of labor and capital for convenience. Prejudice d_i^{ER} is zero for unprejudiced persons and positive otherwise. Finally, we assume throughout that all employers are whites who do not differ in their productive capacity as workers.

Employment of white workers at all employers is determined by the familiar first-order condition, $\frac{\partial f}{\partial L} = w_W$. The first-order condition for black labor, $\frac{\partial f}{\partial L} = w_B + d_i^{ER}$, captures the essence of Becker's insight. Prejudiced employers behave *as if* black workers' monetary wages are higher than they actually are.

Since black and white workers are perfect substitutes in production and utility is linear in discriminatory distaste, each employer simply hires the type of labor that is "cheaper" to him.

An employer therefore hires only white workers if and only if $w_W < w_B + d_i^{ER}$, and hires only black workers otherwise.

Short-run equilibrium

In the short-run, firm size, which we assume to be equal across firms, and the number of firms are both fixed since there is no entry of new potential employers. The distribution of discriminatory tastes among employers is therefore given in the short run. Equilibrium in the short run requires that no employer prefers to hire workers of a different race at market wages, no worker prefers to work at a different firm at market wages, and all workers are employed.³

Becker shows that, in equilibrium, the wage difference between blacks and whites is not determined by the average level of prejudice among all employers, but rather by the prejudice of the most prejudiced employer who hires blacks in equilibrium – the “marginal discriminator”. Efficiency requires that blacks be sorted into higher cost – that is, more prejudiced – firms only after all jobs at less prejudiced firms have been filled. If there are enough employers with $d_i = 0$, all blacks will be employed by these firms. The marginal discriminator is unprejudiced in this case, and there is no equilibrium wage difference between blacks and whites. If, instead, the number of unprejudiced employers is small relative to the number of black workers, full employment requires that at least some blacks must work for prejudiced employers. Wages adjust as blacks are sorted to ever more prejudiced employers until the market clears and until the last person to hire blacks is indifferent between hiring a black and white worker. Equilibrium wages and prejudice of the marginal discriminator, $d_i^{ER*} > 0$, are given by $w_W = w_B + d_i^{ER*}$ such that all workers are employed, and each employer satisfies his demand for labor. Employers with prejudice higher than d_i^{ER*} hire only white workers, and blacks work for other employers. Given a sufficiently smooth distribution of d_i^{ER} across employers, the marginal discriminator is just indifferent between hiring white and black workers.

Long-run equilibrium

³ We describe the equilibrium with inelastic labor supply for simplicity. This assumption is unnecessary.

In the long-run, employers can change size or shut down and new employers can enter the market. An employer chooses capital and labor to maximize utility subject to the constraint that his utility as an employer exceed that from his next best alternative were he to shut down. The original Becker analysis and all subsequent work make the reasonable assumption that this alternative activity is to sell his labor on the labor market and be an employee at another firm. Let the utility from this alternative be V_i^A , and suppose that the employer's wage in that alternative job is w_i^A . If $\tilde{\pi}_i$ represents the money profit associated with the optimal choices of labor and capital, and \tilde{L}_B denotes the number of blacks the employer would choose to hire if he were to stay open, it follows that an employer shuts down his firm if

$$\tilde{\pi}_i - d_i^{ER} \tilde{L}_B \leq V_i^A. \quad (2)$$

The wage received in the alternative job is the same for all (white) employers. To see this, notice that the price at which an employer can sell his labor on the competitive market is simply a wage determined by the marginal product for someone of his type. As employers do not differ in productivity, each employer, were he to become a worker, receives some wage w^A which equals the marginal product for persons of that type.⁴

Importantly, Becker's analysis assumes that the utility in the alternative job depends solely on the wage the person receives in that job. The shut-down decision for an employer is thus

$$\tilde{\pi}_i - d_i^{ER} \tilde{L}_B \leq w^A. \quad (3)$$

A perfectly elastic supply of unprejudiced employers guarantees that $\tilde{\pi} = w^A$, meaning that the shut down decision for an employer simplifies to

$$-d_i^{ER} \tilde{L}_B \leq 0. \quad (4)$$

Expression (4) shows the main implication of the long-run version of the employer discrimination model, and the one that has generated the most controversy. In the long run, free-entry and perfect competition along with constant returns to scale technology imply that no employer with the discriminatory tastes described by Becker can remain in business. It is costly to discriminate,

⁴ An alternative would be to allow for differences in productivity among employers. This skill heterogeneity might be two-dimensional where one dimension, say entrepreneurial skill, makes workers relatively more productive as employers and the other makes them relatively more productive as workers. The resulting comparative advantage might lead to rents which employers could use to "purchase" discrimination. Though we think such a model is realistic, we choose not to introduce such heterogeneity here because we want to show that discrimination can survive competition with free entry of competitors. For a discussion of such a model, see e.g. Heckman (1998).

and the market drives those with high costs out of business. Taste-based discrimination is eradicated by competition.

3. Role Dependent Preferences

In this section we show that a slight modification in what is assumed about preferences and in how we think about the term “employer” changes the results from the employer discrimination model, and suggests a way to model employer and employee discrimination in a unified model rather than with two separate analyses that have been the norm in the preceding literature.

3.a. Clarifying the role of the employer

Recognizing that prejudiced employers earn lower money profits, Becker makes the case that an unprejudiced person could profitably buy the firm from a prejudiced employer at a price equal to current earnings. Both the former prejudiced employer and the new unprejudiced employer are made better off by this transaction. We argue that Becker here characterizes the employer as the owner of the firm’s capital, rather than as the manager who is a labor input. His argument presumes that the prejudiced former employer leaves the market altogether, allowing continued avoidance of blacks. Becker’s argument implies that the return to capital must be equalized across firms, regardless of the racial makeup of the firm or the prejudicial tastes of the owner or employer. We prefer to think of the employer as a residual-claimant manager who rents capital on a competitive market. In this characterization, all employers must pay the same market rental rate for capital.

3.b. The Shut-down Decision Revisited

Given the description of prejudice used by Becker we argue that it is in his role as a supervisory worker at the firm that his prejudice is most relevant. Put this way, the employer’s alternative option, were he to shut down his firm, is to become an employee at another firm. As in the previous literature, we assume that the wage the employer would receive at that alternative firm is w^A for all employers. However, unlike the earlier literature, we do not assume that the employer’s utility in that alternative job is a function only of the wage. Instead, we assume that the prejudice that a person holds as an employer is carried with him if he were to be an employee in another firm. The assumption implicitly maintained in previous analysis is that preferences

are “role-dependent”, and in particular that they are present only when the person plays the role of “employer” and utterly absent when he plays any other role in the labor market. Formally, the variable d_i^{ER} in the previous section denoted the prejudice of an employer i . Now, we stress that d_i^{ER} is person i 's prejudice when he is an employer, and d_i^{EE} is his racial prejudice when he is an employee. Our assumption of portable racial prejudice implies that $d_i^{ER} = d_i^{EE}$. By contrast, the assumption implicitly maintained in Becker's analysis and in subsequent influential criticisms like that of Arrow is that for prejudiced individuals with $d_i^{ER} > 0$, $d_i^{EE} = 0$.

To see the implications of this reasoning, consider a reformulation of the long-run shut down decision presented in the earlier section. Recall, that the standard analysis says that an employer shuts down if

$$\tilde{\pi}_i - d_i^{ER} \tilde{L}_B \leq w^A. \quad (5)$$

If d_i^{ER} and d_i^{EE} represent, respectively, a person's racial prejudice when he is an employer and when he is an employee, and if L_B^A represents the number of black worker's at his likely alternative job, expression (5) is the special case of the condition

$$\tilde{\pi}_i - d_i^{ER} \tilde{L}_B \leq w_i^A - d_i^{EE} L_B^A \quad (6)$$

where $d_i^{EE} = 0$. The way that the previous literature specifies an employer's alternative utility in effect imposes the very strong assumption that a person's racial animus depend only on whether he is an employer; somehow, the same person cast in the role of employee has no animus about the blacks who are his co-workers.

The simplest specification of our assumption of role-independent preferences is to suppose that $d_i^{ER} = d_i^{EE} = d_i$. With this assumption, the long-run shut-down decision becomes

$$\tilde{\pi}_i - d_i (\tilde{L}_B - L_B^A) \leq w^A. \quad (7)$$

Expression (7) shows that the a prejudiced person who happens to be an employer considers how the racial makeup of his incumbent firm compares to that of the firm he would join as employee if he were to shut down and find a job. The model could be generalized to include uncertainty

about the racial make up of the alternative firm.⁵ Notice that, in general, the assumption of role independent preferences dramatically changes conclusions about whether prejudiced employers shut down in the long run.

For example, if the employer expects that the number of blacks co-workers at the job he would get if he shut down is likely to be the same as the number of blacks in his incumbent firm, or $\tilde{L}_B = L_B^A$, he is indifferent between shutting down and remaining open in the long run, even when his costs, net of prejudice, are higher than his competitors'. More generally, since his money income is the same everywhere, a rational prejudiced incumbent employer will only shut down if he expects that his likely alternative job will have *fewer* black workers than he has hired in his firm. Our analysis shows that prejudiced employers do not *necessarily* shut down in the long run, as argued in the previous literature; whether they do or not depends on how racially integrated firms are in the economy overall.

Expression (7) shows that when preferences are role independent, it is the condition $\tilde{L}_B \geq L_B^A$ that determines the shutdown decision. The previous literature reaches the conclusion that discriminators are inevitably driven out of the market by assuming implausibly that $d_i^{EE} = 0$. Our analysis shows that, in fact, discriminators are driven out of the market only if market conditions are such that every rational prejudiced employer can expect to find an alternative job in which fewer blacks are employed as workers. Indeed, when racial animus is portable across roles, prejudiced employers definitely leave the market in the long run only if the market is segregated enough by race to ensure that every prejudiced employer can find job at a firm at which blacks are not employed.

3.c. Unifying the Employer and Employee Discrimination Models

In addition to the employer discrimination model, Becker presents a separate analysis focusing on prejudice among employees, or co-workers. In the standard version of the employee discrimination model, some white workers are presumed to have distaste for interacting with black co-workers. White workers have utility given by

⁵ This uncertainty might be driven, for example, by costly search or by turnover among future co-workers. In either case, the employer bases his expectation of V_i^A on the distribution of possible L_B^A .

$$U_i = w_i - d_i^{EE} L_B \quad (8)$$

where w_i is the wage the worker receives at the job, L_B denotes the total number of black workers with whom he works, and the EE superscript indicates that d is the individual's taste as an employee about his fellow employees. Expression (8) formalizes Becker's notion that a prejudiced employee acts *as if* his wage is lower when he has black co-workers.

The key result from Becker's analysis of the employee discrimination model is that wage discrimination and occupational segregation are distinct and opposing forces. If black and white workers are perfect substitutes in production, firms can be perfectly segregated by race. In this fully segregated equilibrium, no individual must be compensated for working with someone he dislikes. Two separate labor markets form in equilibrium, with workers in each being paid their marginal product. As these are equal, there is no equilibrium wage difference. Wage differences arise in this model to the extent that black and white labor cannot be segregated in equilibrium. Any impediment to segregation may lead to a long-run black-white wage difference even in fully competitive markets. Imperfect substitutability of high- and low-skilled workers with different skill endowments across races is one such natural impediment. In general, however, under the perfect substitution assumption that is assumed in Becker's main analysis, the effect of prejudice held by co-workers is perfect segregation with no wage differences.

The reader will notice that the standard employee discrimination model leads to precisely the same analysis as does the employer discrimination model with portable racial prejudice. This suggests that with this more reasonable assumption of how prejudice operates, it is not useful to analyze the market effect of prejudice in the two distinct models of employer and employee discrimination as has generally been done in the discrimination literature. A worker's utility from working at a given firm is, by our reasoning, simply a function of his wage and of the racial makeup of the firm, and not at all of whether he plays the role of employer or employee at the firm. With role-independent racial preferences, any disutility he receives from cross-racial contact is unaffected by which of these two roles he plays at the firm. And, there is no reason to suppose that his wage is higher or lower simply because he is an employer rather than an employee. The employer at a firm differs from an employee mainly because the employer chooses the particular racial composition of the workers at the firm. No person would accept lower wages for the right to make this choice since he could, in a competitive economy, work at

some other firm at which someone else has chosen the particular racial mix he himself would have chosen.

So long as there is no skill heterogeneity and thus no comparative advantage to being an employer, individuals must, in equilibrium, be indifferent between playing the role of employer or employee at a given firm, and also between playing either role in another firm with exactly the same racial makeup. The wage a prejudiced worker receives as either an employer or employee at firm with a given racial makeup must therefore be lower than the wage he receives in either role at a firm with a higher fraction of black workers as his disutility from cross-racial contact is larger in the latter case.

In sum, the assumption of preferences that are portable across roles leads to a formalization that is indistinguishable from what was called the “employee discrimination” model in Becker’s analysis. Blacks and unprejudiced whites do not care about the racial makeup of their coworkers and are indifferent about the racial makeup of the firms to which they sort themselves as workers. Prejudiced whites by contrast dislike interacting with black co-workers and receive a higher wage for doing so whether as employees or as the particular sort of employee called an employer. The extent to which there is racial discrimination in equilibrium depends exclusively on the ability of the market to segregate prejudiced workers from those they dislike and on the distribution of racial prejudice among whites overall. If perfect racial segregation is costly or otherwise difficult to achieve, some blacks may be forced to work with prejudiced whites, and blacks will receive a lower wage than whites in equilibrium. Efficiency guarantees that cross racial contact will occur between blacks and the least prejudiced whites first, with the most prejudiced whites being the most likely to work in segregated environments. The equilibrium wage gap should be determined not by the overall level of prejudice among whites but rather by the prejudice of a marginal discriminator who, given the number of blacks in the economy and the equilibrium wage, is just indifferent between working within an integrated environment and in a perfectly segregated firm.

3.d. The Example of Imperfect Substitutability

In this section we formally study the effect of imperfect substitutability of labor in production – the most likely reason why the labor market might not be racially segregated in equilibrium. There are other reasons why it might be difficult to segregate the market by race, such as costly

search or the presence of racial employment quotas. We discuss imperfect competition both because the ease of exposition, and because this was an example discussed by Becker in his original discussion of co-worker discrimination.

We suppose that individuals are either black (B) or white (W), and possess either high (H) or low (L) skill. All firms have identical production functions which are constant returns to scale in capital, high-skilled labor and low-skilled labor. We assume that black and white labor of a given skill level are perfect substitutes, but that high and low skilled labor are imperfect substitutes for one another. Constant returns to scale allows us to, without loss of generality, consider a firm to be a single high-skilled worker matched with some number of low-skilled workers. We treat this high-skilled worker as the firm's residual claimant.

In general, a high-skilled worker's utility is

$$U_i^H = \pi_i^H - d_i \times (L_B^L + L_B^H) \quad (9)$$

where $\pi_i^H = f(K, L_B^H + L_W^H, L_B^L + L_W^L) - rK - w_B^L L_B^L - w_W^L L_W^L$, d_i is individual i 's taste for discrimination and L_r^s denotes the total employment of workers from skill-class $s \in (H, L)$ and race $r \in (B, W)$ at individual i 's firm. Since we consider the case of a single high-skilled worker matched with some number of low-skilled workers, this utility function can be re-written

$$U_i^H = f(K, L_B^L + L_W^L) - rK - w_B^L L_B^L - w_W^L L_W^L - d_i L_B^L \quad (10)$$

where the high-skilled labor term is dropped from the production function because we have assumed it to be equal to one.

Short-Run Equilibrium

Firm size is fixed in the short-run. Thus, each high-skilled worker is paired with a fixed number of low-skilled workers. Although the set-up is similar to the short-run version of Becker's employer distaste model, with high-skilled workers serving the role of the employer and low-skilled workers serving the role of employees, we speak instead of high- and low-skilled workers forming teams to stress the fact that there is no structural importance to the role of employer.

The short-run equilibrium is reached as follows. Ordering high-skilled workers by their d_i , low-skilled blacks are allocated first to the least discriminatory high-skilled workers. The last high-skilled worker to be matched with low-skilled black workers is the marginal discriminator. In equilibrium, the black-white wage gap must be just high enough to compel the marginal discriminator to match with low-skilled blacks but not high enough to induce the next most discriminatory high-skilled worker to do the same. All high-skilled workers with d_i greater than that of the marginal discriminator pair with white low-skilled workers; all high-skilled workers with d_i lower than the marginal discriminator's pair with black low-skilled workers. No high-skilled worker prefers to replace his low-skilled workers with those of the other race, and no low-skilled worker prefers to change firms.

If there are not enough non-discriminatory high-skilled workers to match with every low-skilled black, there must be a low skilled racial wage gap in equilibrium, or $w_B^L < w_W^L$. At least one high-skilled prejudiced white must work with low-skilled blacks in this case, and he can be induced to do so rather than work with a white low-skilled workers only if black wages are lower. It is also obvious that there are in equilibrium no firms consisting of a high-skilled worker and both low-skilled blacks and discriminatory low-skilled whites. In such firms, low-skilled whites would have to be compensated for working alongside blacks.⁶

Long-run equilibrium

In the long-run, the size of teams of high-skilled and low-skilled workers can vary. It is useful to speak as if the high-skilled worker chooses the number of low-skilled workers with whom to match. Since labor supply and racial preferences are portable from one job to the next, it is not useful to speak of a firm shutting down. With inelastic labor supply, in equilibrium each worker must work at some firm. A high-skilled worker who shuts down his firm must in equilibrium work at some other firm.

First-order conditions

⁶ It is possible that because of indivisibility and fixed firm size there would be one firm whose low-skilled workers are integrated.

To see the equilibrium matching of high- and low-skilled workers, consider first a non-discriminatory high-skilled worker assessing how many low-skilled workers to team up with. Since black wages are lower than white wages in equilibrium, and since he is indifferent about the race of the low skilled persons in his team, his problem is simply how many black low-skilled colleagues to have at the wage of w_B . The first-order condition for this problem is therefore the familiar

$$\frac{\partial f}{\partial L^L} = w_B^L. \quad (11)$$

Since high-and low-skilled labor are imperfect substitutes, $\frac{\partial f}{\partial L^L}$ is decreasing in L^L . Therefore, given w_B^L there is some optimal level L^{L*} , which would be the optimal low-high skill ratio at every firm if there were no discriminatory tastes. We henceforth refer to L^{L*} as the optimal skill ratio because it is the one that would maximize total physical production in the economy.

Now consider a high-skilled worker with $d_i > 0$. In order to achieve a low-high skill ratio of L^{L*} , he must either suffer the disutility of working with black workers or hire low-skilled whites at the higher white low-skilled wages. Both of these options increase the effective low-skill wage, and thus lower a prejudiced high-skilled workers' demand for all low-skilled workers, relative to the optimal level L^{L*} .

A prejudiced high-skilled worker's first-order condition is

$$\frac{\partial f}{\partial L_B^L} = w_B^L + d_i. \quad (12)$$

If he chooses to work (form a team) with low-skilled blacks, he acts as if the black low-skill wage is higher than its nominal value because of his disutility of working alongside blacks. As a result, he chooses a skill ratio L_B^{L*} that is lower than the optimal skill ratio L^{L*} .

If instead he chooses to work with white low-skilled workers, the first-order condition is

$$\frac{\partial f}{\partial L_W^L} = w_W^L. \quad (13)$$

Since $w_W^L \geq w_B^L$, the optimal L_W^{L*} is also lower than L^{L*} . In equilibrium prejudiced high-skilled workers will form teams with low-skilled blacks if

$$w_B^L + d_i < w_W^L, \quad (14)$$

and with low-skilled whites otherwise.

Equilibrium wage setting

The equilibrium matching of high- and low-skilled workers is as follows. Black low-skilled workers are allocated first to un-prejudiced high-skilled workers. If there are enough un-prejudiced high-skilled workers to match with black low-skilled workers at this ratio then there need not be any wage difference. If however, there are more low-skilled black workers than can be absorbed by non-discriminatory high-skilled workers, some discriminators will be compelled to work with those for whom they harbor distaste.

Prejudiced high-skilled workers are induced to hire blacks by an equilibrium reduction in w_B^L relative to w_W^L . As low-skill wages of blacks fall relative to those of whites, adjustments on two margins serve to clear the market for low-skilled blacks. First, the reduction in w_B^L induces any given un-prejudiced high-skilled person to match with more low-skilled blacks. Second, the reduction in w_B^L relative to w_W^L induces increasingly more prejudiced high-skilled whites to match with low-skilled blacks instead of whites. The wage difference that clears the market for low-skilled blacks and whites is the long-run equilibrium low-skill racial wage gap. There is once again a marginal discriminator who is indifferent between matching with black and white low-skilled workers. All high-skilled workers with d_i greater than the marginal discriminator form firms and work white low-skilled workers; and high-skilled workers with d_i lower than the marginal discriminator's pair with black low-skilled workers. No high-skilled worker prefers to replace his paired low-skilled workers with those of the other race, and no low-skilled worker prefers to change firms.

If the ratio of low- to high-skilled workers is the same across race and all black high-skilled workers are non-discriminatory, the marginal discriminator is black and no racial difference in wages arises. In this case, segregation fully substitutes for market discrimination. Similarly, if the number of blacks relative to whites is very small, then even if there are relatively more low-skilled blacks than whites, the marginal employer of low-skilled blacks is likely to be an unprejudiced white, meaning that there will be no long-run equilibrium wage gap. The situation is different if the number of blacks is relatively large and/or if the ratio of low to high-skilled workers is much higher among blacks than whites. In these cases, the marginal discriminator is likely to be a prejudiced high-skilled white, and there will be a long run racial difference in low-skilled wages. A large number of low-skilled blacks or a relatively large number of high-skilled discriminators means members of these two groups must be matched within firms. As a result there will be a persistent racial wage gap among the low skilled in a competitive market.

How do high-skilled workers fare because of the existence of prejudice? Notice that unprejudiced high-skilled persons, both black and white, benefit from the existence of prejudiced high-skilled whites. The reduction in low-skilled black wages caused by the preferences of these workers leads unambiguously to increased earnings for high-skilled non-discriminators. As residual claimants to revenues from production, the cost savings goes directly to increase their earnings. In addition, the reduction in wages induces them to match with more low-skilled workers, increasing total production at the firm and further increasing their income.

For high-skilled whites who are prejudiced the situation is slightly more complicated. There are two groups of such workers to consider: those who match with low-skill blacks and those who match with low-skill whites. Consider the former group. These are prejudiced persons with $d_i < d^*$ (where d^* denotes the prejudice of the marginal discriminator). Because of their discriminatory tastes, they choose not to take full advantage of the decline in factor costs. The decline in their nominal costs causes an increase in their income, but they do not expand their interaction with low-skilled workers by as much as non-discriminatory whites, and may in fact reduce it. Net of their own discriminatory tastes, these workers may perceive low-skill wages to have increased or decreased as a result of market discrimination. If they perceive low-skill wages, net of their own discriminatory tastes, to have increased and in turn reduce their productive interaction with low-skilled workers, their income may decline.

To see why they may erroneously perceive low-skilled wages to have increased, consider the case where the entire low-skill wage gap is accounted for by a decline in low-skill black wages (i.e. low-skill white wages remain constant). Because the racial wage gap, and by assumption the decline in w_B^L , is equal to d^* , we can be sure that for all whites who match with blacks in equilibrium, $d_i + w_B^L < w_0^L$ (low-skill wages with no discrimination). Net of their discriminatory tastes, these workers perceive net low-skill black wages to have declined as a result of discrimination. Now alternatively consider the case where discrimination causes low-skill white wages to increase. Since we know $w_W^L - w_B^L = d^*$, it must be the case that $w_0^L - w_B^L < d^*$, where the inequality is strict. Any high-skill white with $w_0^L - w_B^L < d_i < d^*$ perceives low-skill wages to have increased.⁷ And the resulting decline in productive interaction with low-skilled workers may overwhelm the effect of paying their low-skilled workers wages below marginal product, leading to a decrease in their own income.

The second group of prejudiced high-skilled workers interacts exclusively with low skilled whites. The wages of this most prejudiced group must fall. As a result of the increase in white low-skill wages, these high-skilled discriminators choose to match with fewer low-skilled workers. It is among this group of workers that we see most clearly Becker's insight that discriminators pay some of the cost of discriminating.

Black and White Net Gains and Losses from Discrimination

The preceding shows that high- and low-skill blacks are affected very differently by discrimination, even though we have assumed white prejudice not to be a function of black skill levels. The market's tendency to segregate discriminators and discriminated tends to separate high-skilled blacks from prejudiced whites. High-skilled blacks benefit from a decrease in the price of a complementary factor of production. This is obviously not the deliberate intention of white racists; it is rather a side effect of their prejudice.

The effect of prejudice is different for different types of whites as well. Only prejudiced whites with the strongest racial feelings lose unambiguously from the existence of prejudice.

⁷ A high-skilled worker who matches with low-skilled blacks will perceive the low-skill wage to have increased net of his own prejudice if $w_B^L + d_i > w_0^L$. We can rewrite this condition $w_0^L - w_B^L < d_i$.

Unprejudiced whites certainly benefit, and whites with mild levels of prejudice may benefit as well.

Is it possible to say something about whether blacks as a group gain more (or lose less) from the presence of racial prejudice in the labor market than do whites? We turn to Becker (1957) to answer this central question. Using the analogy of two countries trading, he reminds us that any reduction in trade must be detrimental to the total income of each of the countries. In our case the countries should be thought of as black and white, and trade as the productive interaction of workers within a racially integrated firm. Because discrimination causes a decrease in the amount of productive interaction across race, each race must lose as a whole from discrimination.

But, who loses more? Becker's analysis shows that the ratio of black to white total income falls as a result of a decrease in trade so long as in a world with no discrimination

$$\frac{I(W)}{I(B)} > \frac{N_B^L}{N_w^L}$$

where I indicates the total income of each race, and N^L indicates total labor supply of low-skilled workers by race. This condition is clearly satisfied if there are fewer low-skilled blacks than low-skilled whites.

4. Empirical Evidence about Racial Prejudice and Wages

4.a. Data description

Remarkably, very little work has been done in economics studying racial prejudice directly, and no empirical work of which we are aware has assessed the connection between prejudice and wages.⁸ Several questions thus remain unanswered. How do racist sentiments vary across regions in the U.S.? Is the level of prejudice to be found among different types of workers – employers versus non-employers – consistent with the long run predictions of the theoretical work of Becker and Arrow? How have racially prejudiced views evolved over time, and have they changed differentially across different types of people or places? Most important for our purposes, what is the empirical relationship between racially prejudiced feelings among whites in

⁸ Two exceptions are Cutler, Glaeser and Vigdor (1999) and Card, Mas, and Rothstein (2006) each of whom relate a subset of the questions studied in this paper to levels of residential racial segregation.

a community and the absolute and relative level of black wages in those places? In particular, is there evidence in the data to suggest that relative black wages are reduced when prejudice is high among whites overall, among white employers, and especially, when the prejudice of the white person who could plausibly be called the marginal discriminator is high?

We use data from multiple waves of the General Social Survey (GSS) [1972-2004] and from multiple waves of the May [1973-1978] and Outgoing Rotation Group [1979-2002] files of the Current Population Survey (CPS) to attempt to answer these questions. The GSS is our source of data on racial prejudice. 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, so it is useful that the various questions posed in the survey over the years touch on the different dimensions along which racist sentiments might manifest themselves. Among other things, respondents were asked over the years such questions as their feelings about interracial marriage, their sense of whether racially restrictive housing covenants were appropriate, their views about children being racially segregated in schools, and their view on whether the government should be obligated to help blacks.

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 prejudice. 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. In much of our analysis to follow, we focus on changes in racial prejudice over time – an exercise which requires identifying a consistent set of preference questions from one year to the next. When doing these analyses we study the six question which *jointly* appear most frequently in the survey. Table 1 lists the GSS variable abbreviation and a summary for each of the full set of 26 racial prejudice questions asked in the GSS. The shaded rows in the table indicate the six questions, used later in the analysis, which jointly appear most frequently in our data.

We use responses from whites aged 18 and older, and recode responses so that higher values correspond to more prejudiced answers.⁹ To permit straightforward aggregation across questions

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

and interpretation of magnitudes, we normalize each variable using the 1977 report. Specifically, we subtract off from responses to each question the mean of the response to that question in 1977, and divide by the standard deviation of answers measured in the first year the question was asked.¹⁰ Two of the variables (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. For questions that were introduced to the survey subsequent to 1977 or that were asked only in 1972, we normalize by both the mean and standard deviation measured the first year the question was asked.

4.b. Definitions of various prejudice indices

Rather than present results separately for each of the 26 questions, we form a uni-dimensional index of prejudice. To compute this index, we compute the average of the normalized responses described above to the full set of prejudice questions, separately for each individual*year observation in the GSS. Call this variable \bar{d}_i ; call the mean taken over the six questions referenced above \bar{d}_i^{subset} . Our empirical work focuses on community level prejudice, where the communities we focus on are the nine census divisions. These are the smallest areas of aggregation available in both the GSS and CPS. We take various averages of the individual measures \bar{d}_i to compute these community prejudice measures. Our analysis uses two main measures for three subgroups of the population: all whites, white supervisors (defined based on occupation), and high-skilled whites (as indicated by those with at least 16 years of education). The first measure of prejudice is the average of \bar{d}_i by census division (indexed by r) over all individual*year observations within the subgroup. We denote this measure \bar{d}_r for all whites, \bar{d}_r^S for white supervisors, and \bar{d}_r^H for high-skilled whites.

We present results using measures of prejudice for high-skilled whites for two reasons. First, the imperfect substitution model discussed above implies that it is the relevant measure of prejudice

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.

in determining equilibrium black-white wage differences. Second, being a college graduate is an alternative indicator (relative to being self-employed) of the kind of person likely to be an “employer” in the Becker model.

Our second measure of the prejudice in a community is meant to approximate as closely as possible the theoretical construct of the marginal discriminator – that person whose racial views determine the equilibrium black wage gap if prejudice is indeed an important part of the explanation for equilibrium discrimination. The logic we use to construct this measure is straightforward and flows naturally from the reasoning presented earlier in the theoretical discussion. Suppose that p percent of the population is black. If racial prejudice governs hiring and compensation in the manner earlier described, then efficiency requires that workers will be sorted first to the least prejudiced employers. If we assume that employers are essentially the same with respect to firm size, then it follows that blacks will work at the p percent least prejudiced of all employers. A natural measure for the prejudice of the marginal discriminator, \hat{d}_r^{S*} , is thus the p^{th} -percentile of the census-division-specific distribution of \bar{d}_i^{subset} among white supervisors, where p is the share of the census division that is black over the full CPS sample period.¹¹ Motivated by the idea that employer status is endogenous, we also compute this “marginal discriminator” measure of prejudice for all whites, denoted \hat{d}_r^* , and for high-skilled whites, denoted \hat{d}_r^{H*} .

Notice also that none of the community prejudice measures described above varies over time. In some of our analysis, we wish to examine the effect of *changes* in a community’s prejudice over time. To compute a community-wide time-varying measure \bar{d}_{rt} , we take the average of \bar{d}_i^{subset} by census-division*year. As we will show below, virtually all of the variation in prejudice is across census divisions, and not within areas over time. We therefore focus our empirical tests on the census-division level indices that are pooled over the 30 year sample of the GSS.

¹¹ We construct our measures of the marginal discriminator’s prejudice as the p^{th} percentile of the distribution of \bar{d}_i^{subset} and not \bar{d}_i because it is based on a constant number of questions in each year. Since \bar{d}_i is an average over responses to different numbers of questions in different years, the variance of \bar{d}_i will tend to be higher in years in which a smaller number of questions happened to be asked. Statistics based on the tails of the distribution of \bar{d}_i would disproportionately measure prejudice in those years in which a relatively small number of questions were asked.

4.c. Patterns and trends in prejudice in the GSS

We begin by describing the general patterns of prejudice across regions of the U.S. and over time. Table 2 shows six of the prejudice indices for each of the nine census divisions. As described above, higher values indicate greater racial prejudice. All six region-level prejudice measures are very highly correlated. Each index shows that 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 (Alabama, Kentucky, Mississippi, Tennessee), and next greatest in the South Atlantic (Delaware, District of Columbia, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, West Virginia), and West South Central (Arkansas, Louisiana, Oklahoma, Texas) divisions. Prejudice is least severe in New England (Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, Vermont), and in the Pacific (Alaska, California, Hawaii, Oregon, Washington) and Mountain (Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, Wyoming) divisions. We suspect these patterns match many readers' priors.

The magnitude of differences in measured prejudice across regions is also noteworthy. The difference in \bar{d}_r between the East South Central and New England census divisions is on the order of one-half of an *individual-level* standard deviation. To put this difference in perspective, the median East South Central respondent has the same \bar{d}_i as the 82nd percentile respondent from New England. The median-prejudiced New England respondent would be at the 22nd percentile of the East South Central prejudice distribution.

The first three columns of the table list alternative indices of prejudice, each based on a single question from GSS. Each of the three questions have binary answers and were asked in a large number of years. As is clear in the table, the correlation of the black-white wage gap with the indices based on the full set of questions is not driven by the constant weights we have chosen across questions. Census divisions where a larger fraction supports a law against interracial marriage have significantly larger black-white wage gaps. The same is true of census divisions where more would not vote for a black candidate for President. Interestingly, the correlation is weaker when prejudice is measured according to the fraction that oppose busing to integrate the

schools. In a sense, this pattern is reassuring since we think the former two questions do a better job of getting at the type of prejudice that would cause someone to dislike working alongside blacks. In contrast, the latter question may pick up feelings about the role of government in individual affairs. Rather than make such ad hoc judgments, we take a hands-off approach in the formal analysis below and weight each question equally.

It is interesting to notice that blacks live most predominantly in regions of the country where racial prejudice is most severe. This pattern may indicate that prejudice is caused in part by cross-racial contact and by competition for economic resources. Furthermore, as with any survey data, significant care must be taken in interpreting these responses. Another concern is that regional differences in answers to the GSS questions indicate differences in candor about true racial feelings rather than differences in actual prejudice. If this is the case, then we should find no relationship between measured prejudice and wage differences. A more serious concern for our purposes is that the indices are correlated with unobserved regional differences in productivity between black and white workers. This is a concern we address below.

Figure 1 shows trends in responses to the twelve most commonly asked GSS prejudice questions, averaged across the entire sample of whites. The figure reveals a general downward trend in reported racial prejudice. 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. Declines range from close to zero for questions concerning government treatment of and spending for blacks, to more than a half a standard deviation for questions concerning laws about school and residential integration. In the midrange of declines, at about a quarter to a third of a standard deviation, are questions about the willingness to vote for an otherwise qualified black presidential candidate, support for laws concerning interracial marriage, and the desired racial composition of schools.

Figure 2 shows trends in the uni-dimensional prejudice measure \bar{d}_r plotted separately for each census division. Three key things should be noted about the figure. First, the decline in measured prejudice has been widespread. Between 1977 and 1996 (the years over which \bar{d}_r is available) measured prejudice declined in each of the nine census divisions. Second, the relative rankings of prejudice across regions has been preserved. New England, and the Pacific and

Mountain divisions remain the least prejudiced, while the East South Central and South Atlantic divisions remain the most prejudiced, according to our index. Third, there has been remarkably little convergence across census divisions in the level of measured prejudice. As a result, virtually all of the variation in prejudice in our data is across areas and not within areas over time.

In short, there is a great deal of variation across regions of the U.S. in levels of measured prejudice. Racial prejudice appears to be greater in the South and lesser in the Northeast and West. Furthermore, there has been a marked decline in measured prejudice over the past 30 years. This variation in tastes raises a natural question: are black-white wage gaps largest in the places with the most severe prejudice? The answer appears to be yes, and the correlations are striking. Table 2 shows that black wages are lowest relative to whites in precisely the census divisions where measured prejudice is greatest. Figure 3 shows this relationship clearly. With only 9 observations, the point estimate of -0.186 is statistically significant ($t = 3.44$) and the univariate regression has an R^2 of 0.628.

4.d. The prejudice of supervisors and the self-employed

Before turning to a formal analysis of the empirical relationship between area-level prejudice measures and the black-white wage gap, we examine a pair of implications of the model described above. We emphasized above Becker's insight that the market pressures prejudiced workers to seek jobs in firms with few blacks. It is their ability to do so—the market's ability to segregate—that determines whether racial wage gaps persist. This insight gives rise to a testable implication of the model: that more prejudiced workers should work in firms with fewer blacks. Ideally, we would conduct this test using firm-level data with both measures of prejudice and racial composition. In the absence of these data, we investigate whether more prejudiced workers are more likely to become self-employed. We know from other sources, that the vast majority of self-employed have no employees. Self-employment is therefore an effective way to avoid working with blacks.

The first two columns of table 3 shows comparisons of the responses to each of the prejudice questions by white self-employed and non-self-employed. Two regression coefficients are reported, each one from a separate regression. The left-hand coefficient is the unconditional difference in \bar{d}_i between self-employed and all other white workers. A positive point estimate

indicates that employers are more prejudiced. The right-hand coefficient is the difference estimated in a regression with education, age, region, and year dummies.

Consistent with the predictions of the model, the results show the self-employed are statistically significantly *more* prejudiced than the average worker. This is true both conditionally and unconditionally. The smaller point estimate in the specification with controls suggests that the difference is smaller among those with the education and experience levels typical of the self-employed. While the point estimate is small, we take this as evidence consistent with the idea that more prejudiced workers seek out jobs separated from those they dislike.

We contrast this result with a comparison of the average prejudice of supervisors and non-supervisors. Because the role of supervisor is endogenously chosen, we should expect supervisors to have the same tastes on average as non-supervisors who have the skills necessary to be supervisors. This is contrasted with the traditional interpretation of Becker, which says only the unprejudiced should remain supervisors, thereby predicting that supervisors should be significantly less prejudiced than the average worker. These comparisons, shown in columns 3 and 4 of table 3, show that white supervisors are unconditionally more prejudiced than all other whites. However, conditioning on age, education, year and census-division dummies reduces the difference to zero.

4.e. The relationship between measured prejudice and the black-white wage gap

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 1973 to 1978 with the Merged Outgoing Rotation Group (MORG) files from 1979 to 2002.¹² The sample includes full-time black and white males aged 16 to 64.¹³ Our basic specification could be estimated in two ways. We could estimate an OLS regression of log wages on education, a quadratic in potential experience, race-specific year dummies, a black dummy, \bar{d}_r , and the interaction of \bar{d}_r and the black dummy. The coefficient of interest would be the

¹² 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.

¹³ Results are similar using a sample of all black and white males aged 16 to 64 with positive earnings.

estimate on the interaction, and negative values would tell us how much lower black wages are relative to whites in census divisions with greater measured prejudice. Because the index of prejudice only varies at the census-division level, we are concerned this procedure might underestimate standard errors, even if we were to cluster at the census-division level.

Instead, we estimate the model in two steps. First, we estimate the residual black-white wage gap in each census division. We do so by estimating an OLS regression on the variables described above, leaving out the prejudice index and including a separate black dummy for each census division. The estimated effects on each of these nine black dummies become the dependent variable in the second step, in which one or more of the prejudice indices are independent variables of interest. The second step regression has nine observations and therefore produces conservative standard error estimates. The results are substantively almost identical to those from the one-step procedure described above.

The results are presented in table 4. The table is broken into three panels. We focus on the top panel first. The basic result is shown in column 1 of the top panel. The result shows what we see in figure 3, that more prejudiced census divisions have larger black-white wage gaps. The standard deviation of the prejudice indices are about 0.6. The result therefore implies that a one-standard deviation increase in average prejudice is associated with black wages that are 0.097 log points lower. The result is estimated precisely enough to reject a zero correlation at any reasonable significance level. Inspired by Becker's insight, column 2 reports the estimated relationship between the racial wage gap and the index of the marginal discriminator. As described above, the marginal prejudice index is the p^{th} -percentile of the individual-level prejudice distribution, where p is the percent black in the census division. It is also the case that census divisions with more prejudiced "marginal" discriminators have lower relative black wages.

Becker's model makes a stronger prediction, however. Becker points out that because the market pressures discriminatory workers and blacks to sort away from each other, it is not the average level of prejudice that determines the racial wage gap in equilibrium. It is the level of prejudice of the marginal discriminator. Therefore, conditional on the average level of prejudice in the labor market, the marginal index should be significantly related to wage gaps. Indeed, this is the case. In column 3, we report a specification that includes both the average and marginal indices and the coefficient on the marginal index is large and significant. In column 4, we report one

further falsification exercise aimed at testing whether the black-white wage gap is indeed related to the level of prejudice in the left-tail of the prejudice distribution and not the right-tail. We create a false marginal index, which we call the “mirror”, by taking the 100- the p^{th} -percentile of the individual-level prejudice distribution. In a sense this is the mirror image of the marginal index. The specification reported in column 4 includes the average, marginal and mirror indices and only the marginal index has a significant effect. In fact, each of the other point estimates is positive and close to zero. We take this as strong evidence in favor of Becker’s model.

A valid concern one might raise after viewing the preceding results is that regions with more severe measured prejudice also have other unobserved characteristics that negatively affect black wages more than white wages. Though we control for the most obvious suspects, such as differences in education levels, there are always other possibilities. One example is that school quality (not quantity) may be relatively worse for blacks in places with more prejudice.¹⁴ In fact, this may be a direct result of prejudice as in the case of segregated schools. If this were the case, we could see the negative relationships observed in the data even if there were no direct effect of prejudice on wages in the labor market.

The specifications that include both the average and marginal indices provide a potential solution to this problem. In each of these specifications, we include a control for the average level of prejudice in the census division. Becker pointed out that these variables should not affect relative wages of blacks and whites in a direct way. They should, however, be correlated with the unobservables just described (for instance, because they represent the prejudice of the median voter). The results in columns 2-4 of table 4 show that black relative wages are related to the prejudice of Becker’s marginal discriminator, *even conditional on average prejudice levels in the region*.

The bottom two panels of table 4 report the same specifications but for prejudice indices computed among the sample of high-skilled (college graduates) and supervisors. We take these to be imperfect measures of “employer prejudice”. The basic pattern of results is remarkably similar for these specifications. The black-white wage gap is significantly related to the level of prejudice, both average and marginal, among high-skilled and supervisors. In both sets of results,

¹⁴ Notice that it would not cause a bias if school quality were relatively lower for both blacks and whites since we include region effects, unless the effect of school quality on wages differed by race.

the relationship is strongest with the marginal index, though the estimates are only marginally significant.

5. Conclusion

Our goal in this paper has been to argue that the null under which economists have recently operated, that employer prejudice is not an important part of the explanation for observed racial wage gaps, might be incorrect. Two lines of argument have been adduced in support of our claim that racial prejudice of the sort first formalized by Becker (1957) in his seminal analysis of labor market discrimination may indeed matter for observed minority wages.

The first part of our analysis re-evaluated the theoretical underpinning of the null mentioned above. We show that under the reasonable assumption that racial prejudice is portable across roles – that a prejudiced employer in one firm would likely be a prejudiced employee at another – it is not necessarily the case that prejudiced employers will be driven out of the market in the long-run. Instead, we show that if it is costly to separate the market by race, a prejudiced employer can very well remain in business. Using the specific example of imperfect substitution, we show that with portable preferences racial wage gaps arising from employer discrimination can theoretically last into the long-run, with the size of the wage gap determined by the prejudice of the least prejudiced person to hire black workers.

The second part of our analysis is an empirical examination of racial prejudice and of the connection between prejudice and wage gaps. Strangely, in the large previous literature on discrimination in economics, we have been unable to find any previous work directly studying reported racial prejudice. 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 varied widely across regions.

We present several pieces of empirical evidence about reported prejudice which suggest that, consistent with our theoretical argument, prejudice might indeed be an important source of racial wage differences. We first show that self-employed whites are more prejudiced on average than

other whites. We take this to be evidence in support of the idea that the market pressures prejudiced workers to seek jobs at firms with little or no exposure to blacks. We contrast this result with the fact that conditional on various human capital controls supervisors are just as prejudiced as the average worker. This result is inconsistent with the prediction that prejudiced employers should be driven from the market.

Next, using individual-level data from the Current Population Survey, we directly examine the relationship between the racial wage gap and the level of prejudice in a community. We show that the racial wage gap is larger the higher the level of overall prejudice in a community. We find the same basic patterns for alternative measures of community racial prejudice, but the strongest effects are for results in which we measure community prejudice using our estimate of the prejudice of the marginal discriminator. This result is striking because the marginal discriminator's prejudice is shown to be so strongly related to the relative wage gap even in regressions which control for the average level of prejudice among employers overall. More importantly, that the marginal discriminator's prejudice seems so important for wages is precisely what Becker's original analysis as well as our treatment with portable preferences predict.

In our view, the paper's various results, both theoretical and empirical, point to a larger role for racial prejudice in wage determination for minorities than has been acknowledged in the recent discrimination literature. 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. Similarly, while the empirical evidence we have presented is strongly suggestive of an important role for racial prejudice, we have been careful to stress that absent quasi-experimental evidence, causal interpretations cannot necessarily be given to these 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 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.

Brown, Charles (1984). "Black-White Earnings Ratios Since the Civil Rights Act of 1964: The Importance of Labor Market Dropouts." Quarterly Journal of Economics 99 (February), pp. 31-44.

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. Amersterdam: 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 Levit (2004). "The Causes and Consequences of Distinctively Black Names," Quarterly Journal of Economics, 119(3), pp 767-805.

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.

Hellerstein, Judith K., David Neumark, and Kenneth R. Troske (1997). "Market Forces and Sex Discrimination." NBER Working Paper 6312.

Jovanovic, Boyan (1979). "Job Matching and the Theory of Turnover" Journal of Political Economy 87, pp. 972-990.

Katz, Lawrence F. and Kevin M. Murphy (1992). "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." Quarterly Journal of Economics. 107(1), pp35-78.

Lang, Kevin (1986) "A Language Theory of Discrimination", Quarterly Journal of Economics, 101 (May), pp. 363-82

Loury, Glenn C. (1977). "A Dynamic Theory of Racial Income Differences," in P.A. Wallace and A.M. LaMond, eds., Women, Minorities, and Employment Discrimination. Lexington, MA: D.C. Heath and Co.

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.

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.
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.
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 6 questions shaded in gray were asked in the 1977, 1985, 1988, 1989, 1990, 1991, 1993, 1994 and 1996 waves of the GSS. We use these six questions to construct the prejudice indices that vary within region over time, as well as the indices of the marginal discriminator's 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. No question was asked in every year and some were asked in only one.

Table 2: Means of prejudice indices by census division

	Support Law Against Interracial Marriage	Would not vote for Black for President	Oppose busing to integrate schools	Prejudice Indices								% Black	$\overline{\log w_B - \log w_w}$
				Overall		High-skilled		Supervisors					
				Average	Marginal	Average	Marginal	Average	Marginal				
E. Sou. Central	0.435	0.265	0.752	0.141	-0.995	-0.054	-0.728	0.174	-0.630	14.2	-0.231		
South Atlantic	0.321	0.199	0.748	0.004	-1.041	-0.284	-1.045	-0.044	-0.773	16.9	-0.185		
W. Sou. Central	0.273	0.183	0.759	-0.052	-1.084	-0.219	-1.084	-0.076	-1.084	9.7	-0.188		
E. Nor. Central	0.220	0.134	0.766	-0.142	-1.216	-0.407	-1.216	-0.178	-1.123	6.9	-0.129		
W. Nor. Central	0.226	0.145	0.739	-0.181	-1.385	-0.423	-1.385	-0.248	-1.385	2.2	-0.104		
Middle Atlantic	0.178	0.119	0.718	-0.196	-1.159	-0.427	-1.216	-0.215	-1.084	8.6	-0.160		
Mountain	0.152	0.100	0.714	-0.300	-1.385	-0.436	-1.385	-0.287	-1.385	1.7	-0.143		
Pacific	0.119	0.092	0.727	-0.319	-1.348	-0.480	-1.385	-0.339	-1.216	4.5	-0.138		
New England	0.144	0.082	0.694	-0.363	-1.385	-0.555	-1.385	-0.356	-1.385	2.4	-0.134		
Total	0.230	0.146	0.735	-0.156	-1.222	-0.365	-1.203	-0.174	-1.119	7.7	-0.157		
Dep. Var.: B-W Wage Gap, Bivariate OLS Coefficients	-0.305 (0.089)	-0.526 (0.144)	-0.585 (0.579)										
Observations	9	9	9										
R-squared	0.63	0.66	0.13										

Table 3: The relative prejudice of self-employed and supervisors

	Dependent Variable: \bar{d}_i			
	Self-employed		Supervisors	
	(1)	(2)	(3)	(4)
	0.043 (0.010)	0.028 (0.009)	-0.037 (0.010)	0.008 (0.009)
Controls:	N	Y	N	Y
R ²	0.001	0.162	0.000	0.166
No. Obs.	35,950	35,781	38,306	38,111

Note: Dependent variable (\bar{d}_i) is the individual-observation-level average of the 26 GSS prejudice questions listed in table 1. Higher values of \bar{d}_i correspond to greater racial prejudice. The reported estimate is the coefficient on a self-reported indicator for being self-employed in columns 1 and 2 and on an indicator for supervisors, which is based on occupation, in columns 3 and 4. Controls included in the specification reported in column (2) are 20 education dummies, 71 age dummies, 24 year dummies, and 8 census division dummies. Standard errors are clustered by region (census division).

Table 4:
The relationship between black relative hourly wages
and area-level measures of prejudice

<i>Overall:</i>				
	(1)	(2)	(3)	(4)
Average	-0.161 (0.065)		0.082 (0.125)	0.059 (0.143)
Marginal		-0.201 (0.053)	-0.278 (0.129)	-0.250 (0.151)
Mirror				0.016 (0.035)
Observations	9	9	9	9
R-squared	0.46	0.68	0.70	0.71
<i>High-skilled:</i>				
	(1)	(2)	(3)	(4)
Average	-0.200 (0.059)		0.041 (0.164)	0.045 (0.180)
Marginal		-0.149 (0.034)	-0.176 (0.113)	-0.176 (0.123)
Mirror				0.004 (0.027)
Observations	9	9	9	9
R-squared	0.62	0.73	0.73	0.73
<i>Supervisors:</i>				
	(1)	(2)	(3)	(4)
Average	-0.179 (0.055)		-0.026 (0.126)	0.022 (0.151)
Marginal		-0.120 (0.030)	-0.105 (0.079)	-0.149 (0.106)
Mirror				-0.029 (0.044)
Observations	9	9	9	9
R-squared	0.60	0.69	0.69	0.72

Note: Each of the three panels reports results from four regressions. The dependent variable is the census-division-level residual black-white wage gap conditional on race-specific year effects, education, a quadratic in experience, and census division fixed effects. The results reported in the table are from a second step regression with nine observations, one for each census division. Data are from the May (73-78) and Outgoing Rotation Group (79-02) files of the CPS. Sample includes black and white men, age 16-64, who worked at least 35 hours last week. The average measure of prejudice is an index of responses by all whites to questions about racial prejudice from the General Social Survey (GSS). The specific questions are listed in table 1. The index is pooled over time for the full GSS sample period, separately for each of the nine census divisions. It implicitly weights questions by the number of years they are asked in the GSS. The marginal prejudice index is defined as the p^{th} -percentile of the distribution of responses by all whites to a subset of six questions that are asked in each of the following years: 1977, 1985, 1988, 1989, 1990, 1991, 1993, 1994, and 1996, where p is the fraction black in the census division. The mirror, or false marginal, is the 100- p^{th} percentile of the same distribution and is meant to measure variation in the right-tail of the prejudice distribution.

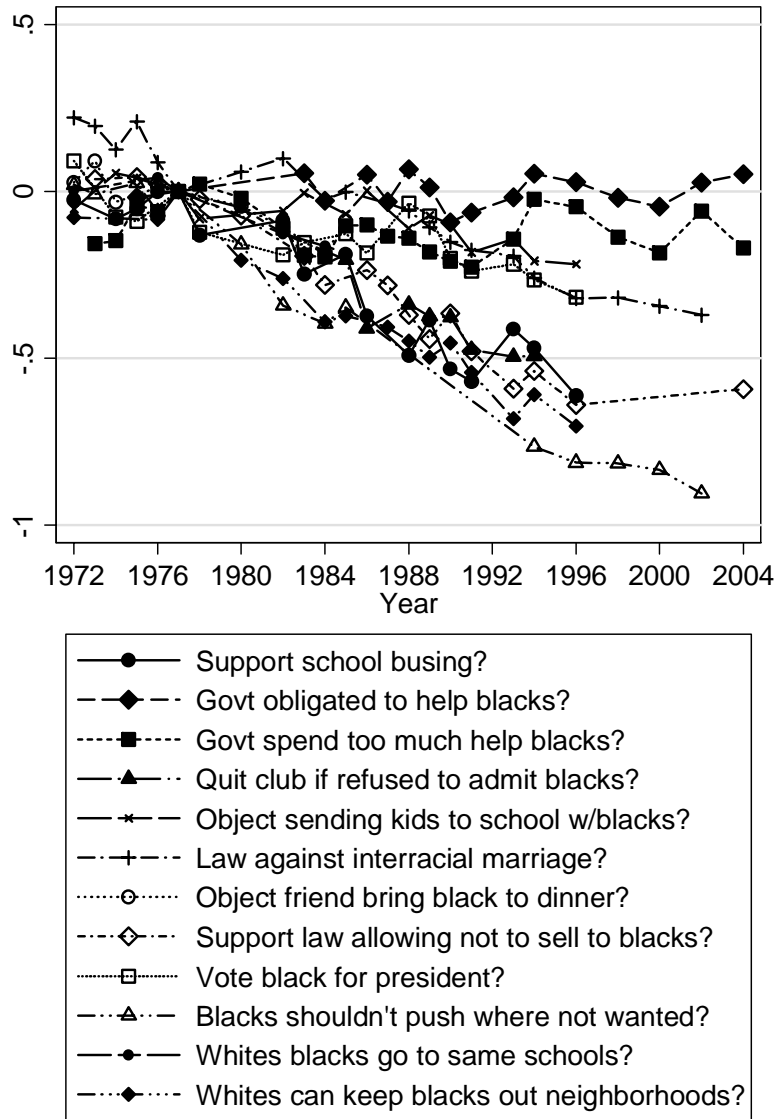


Figure 1:
Trends in responses to GSS prejudice questions

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




Figure 2: Trends in prejudice by census division

Note: Figure plots \bar{d}_{rt} over time by census division for each of the years for which it is available.

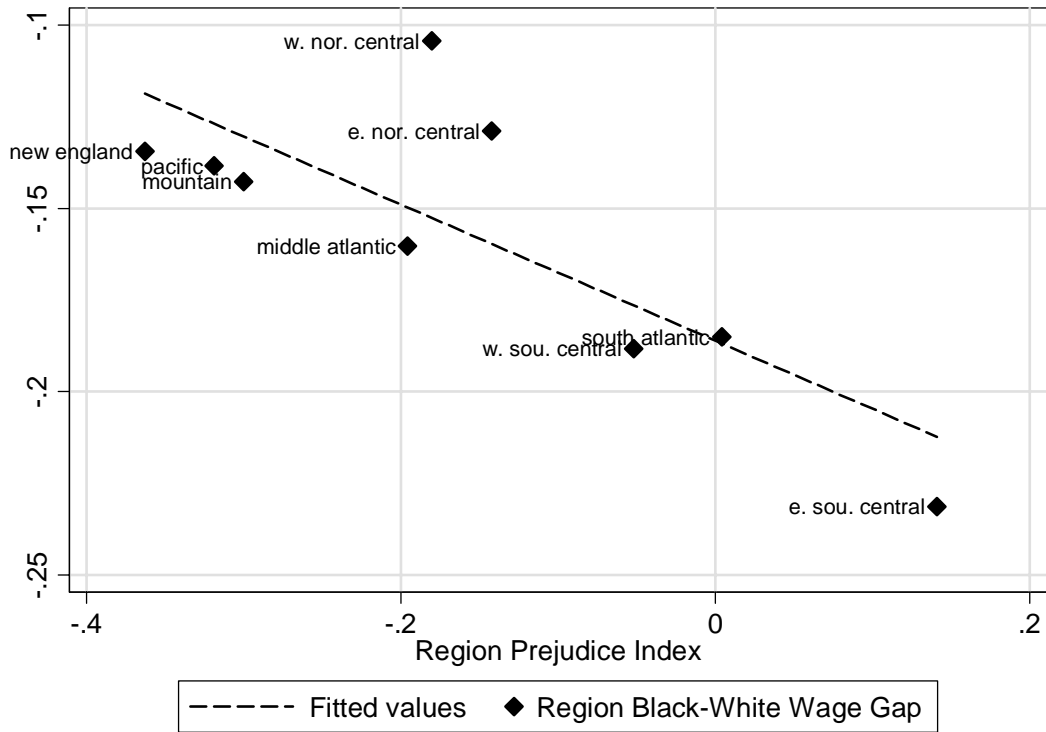


Figure 3:
 Plot of the regional black-white wage gap against an index of regional prejudice (\bar{d}_r)

Note: The estimated slope is -0.186 (standard error, 0.054). $R^2 = 0.628, n = 9$.