Racial Bias in the Manager-Employee Relationship:

An Analysis of Quits, Dismissals, and Promotions at a Large Retail Firm

Laura Giuliano
David I. Levine
Jonathan Leonard

Abstract: Using data from a large U.S. retail firm, we examine how racial matches between managers and their employees affect rates of employee quits, dismissals, and promotions. We exploit changes in management at hundreds of stores to estimate hazard models with store fixed effects that control for all unobserved differences across store locations. We find a general pattern of own-race bias in that employees usually have better outcomes when they are the same race as their manager. But we do find anomalies in this pattern, particularly when the manager-employee match violates traditional racial hierarchies (for example, nonwhites managing whites).

Laura Giuliano is an assistant professor of economics at the University of Miami. David I. Levine is a professor of economics at the Haas School of Business, University of California, Berkeley. Jonathan Leonard is a professor of economics at the Haas School of Business, University of California, Berkeley. The authors are grateful for funding from the Russell Sage Foundation and from the University of California Institute for Research on Labor and Employment, and also for the generosity of the studied employer in sharing their data and time. For helpful comments and suggestions, they thank David Card, Jed DeVaro, Gary Fields, David Lee, Phil Robins, Quincy Stewart, three anonymous referees, and seminar participants at Cornell University, the University of California-Berkeley, the University of Miami, and Florida International University. The data used for this article is proprietary and cannot be made public.

I. Introduction

Does own-race bias in the manager-employee relationship affect employment outcomes? While evidence of own-race bias has been found in several contexts, studies of such bias in the employment relationship have focused on hiring outcomes. The present paper extends the literature by looking at whether own-race bias affects post-hire outcomes. Using personnel records from a large national U.S. retail firm, we examine how racial matches between managers and their employees affect employee rates of quits, dismissals, and promotions.

We exploit changes in management at hundreds of stores to estimate hazard models with store fixed effects that control for all unobserved differences across store locations.[1] We find a general pattern of own-race bias in that employees usually have better outcomes when they are the same race as their manager. But we do find anomalies in this pattern, particularly when the manager-employee match violates traditional racial hierarchies.

In quits, we find evidence of modest own-race bias for whites, Hispanics, and Asians, though not for blacks. The relative quit rates of white, Hispanic, and Asian employees are six percent lower under own-race managers. However, further analysis suggests bias may be strongest among white employees. We find that for whites alone, having an own-race manager significantly reduces the probability that "dissatisfaction with supervisor" is given as the reason for quitting. Moreover, we find evidence that white job-seekers self-select into workplaces based on preferences for own-race managers, and hence that the own-race effect on quit rates may understate the level of own-race bias among white workers.

In dismissals and promotions, we find a common pattern. On one hand, black, Hispanic, and Asian employees have lower relative rates of dismissal and higher relative rates of promotion when their manager is the same race. The evidence of own-race bias is particularly strong for blacks; the relative dismissal rate of blacks is 19 percent lower under black managers than non-black managers, and the relative promotion rate is 79 percent higher. On the other hand, the pattern of bias is the opposite for white employees. Under own-race managers, whites have relative outcomes that are similar or even *worse* than under other-race mangers.

Our main results are generally consistent with theories that assume own-race bias. We find own-race bias in nine of the twelve cases examined (three outcomes, four race groups). But we also find anomalies in this pattern, and interestingly these anomalies are all consistent with status and social identity theories. Such theories suggest that racial differences may evoke different responses in norm-breaking relationships (minorities managing whites, for example) than in relationships that conform to traditional hierarchies. These theories could thus explain two of our most striking findings. First, when white employees work under other-race managers, they have similar or even *better* outcomes than when they work under white managers. But second, despite this preferential treatment, white employees still appear to harbor the strongest bias against other-race managers.

II. Previous Literature on Own-Race Bias—Theories and Existing Evidence

A. Theoretical Background

Economists have tended to explain labor market discrimination in terms of either statistical discrimination or in-group biases. In statistical discrimination (Arrow 1973), a firm or individual uses the average characteristics of a group to make decisions about an individual from that group. Such discrimination does not obviously explain the pattern of own-race bias in our results. It should yield a uniform tendency by all groups to discriminate against a given group, but not a differential tendency to discriminate correlated with own-race matches.

Theories of in-group bias fall into two categories: tasted-based and efficiency-based.

Taste-based models (Becker 1972) posit that people prefer to interact with members of their own race, and that individuals are willing to pay a cost to indulge their preferences. Taste-based models thus predict that workplaces will tend to be racially segregated, and they maintain that discrimination is costly. Efficiency-based models also predict workplace segregation, but posit own-race biases that are driven by efficiency considerations. Employers may prefer same-race matches because of efficiency gains that result either from racially correlated preferences for workplace collective goods (Epstein 1992) or from reduced costs of communication and mentoring (Lang 1986; Athey, Avery, and Zemsky 2000).[2]

In-group theories have usually been applied to the hiring process, and one might ask whether they also apply to post-hire outcomes. Why would managers and employees who would discriminate *ex post* choose to work together in the first place? For biased managers, it may be that they face strong legal and social pressure to hire a workforce that reflects the racial composition of the applicant pool. But once employees are hired, these managers may require higher productivity thresholds from other-race employees. Or they may find that other-race employees are less productive due to poorer communication or less supportive mentoring.[3]

Though biased workers have no legal pressure to accept jobs with other-race managers, they might do so depending on the weight they give compatibility with their manager in the utility function. And if such compatibility were to become more salient after they are hired, then discrimination in quit rates could occur. Further, because employees typically play no role in the hiring of new managers, quitting would be the main way employees who receive new managers could exercise discriminatory preferences. This suggests that employee discrimination in quits should be most prevalent among those who did not “choose” their managers—a prediction we test by looking at employees who get new managers.

More recently, economists have also begun to consider the relationship between discrimination and theories of status and social identity (for example, Akerlof and Kranton 2000; Ridgeway 2007). Both taste-based and efficiency-based models of in-group bias ignore the effects of differing social contexts. However, theories of status and identity suggest the effect of in-group bias in the workplace must be considered in conjunction with whether workplace roles conform to socially accepted roles and hierarchies. In particular, both status and identity theories suggest that racial differences may evoke different responses in norm-breaking relationships than in traditional relationships.

For example, white employees with non-white managers may feel such relationships are a threat to their status or identity. Hence, because they are subject both to own-race bias and to an identity threat, whites might be more likely than other employees either to avoid working for a different-race manager or to quit when they have a different race-manager. On the other side, non-white managers may know that white employees are less willing to accept their authority and are more likely to quit. Also, non-white managers who exercise authority over whites may find that such norm-breaking causes anxiety and psychological discomfort. Hence even if non-white managers hold own-race biases, such factors might still lead them to be relatively deferential toward whites and so dismiss whites less often and promote them more often.

B. Previous Empirical Literature

Evidence of own-race bias has been found in a variety of relationships and outcomes. In superior-subordinate relationships, the economics literature has found own-race bias in outcomes such as arrest rates (Donohue and Levitt 2001); vehicle search rates by police officers (Antonovics and Knight 2004); evaluations of students by teachers (Dee 2005); and foul-calling by NBA referees (Price and Wolfers 2007). However, a recent study of racial bias by police officers who issue speeding tickets finds no evidence of own-race bias; rather, it finds that black and Hispanic officers are relatively harsh on minority motorists and lenient toward white motorists (Anbarci and Lee 2008).

Studies of own-race bias in the manager-employee relationship are found in several literatures: psychology, organizational behavior, and economics. In psychology and organizational behavior, the most germane findings are that subordinates with own-race managers have higher job satisfaction (Wesolowski and Mossholder 1997), more supportive mentoring relationships (Thomas 1990), and higher performance ratings (Stauffer and Buckley 2005). In economics, the research has focused on hiring.[4] In particular, Stoll, Raphael, and Holzer (2004) analyze data from four large cities and find that firms with black hiring agents hire more blacks than those with white hiring agents. Using the same data employed in the present study, Giuliano, Leonard, and Levine (2009) find evidence of own-race bias in hiring for whites and blacks, and also for Hispanics in locations with large Hispanic populations.

The present paper extends the literature on own-race bias in several ways. First, we look at the effects of manager-employee similarity on post-hire employment outcomes. Second, by decomposing exits into manager decisions (dismissals) and employee decisions (quits), we can identify both managerial bias and employee bias. Prior studies have not identified bias in employees or subordinates because they looked only at the decisions of superiors (for example, punishments by police officers) or at outcomes that conflated the decisions of both superiors and subordinates (for example, hiring).

Third, our findings help interpret the results of the hiring studies by Stoll et al. (2004) and Giuliano el al. (2009). Our main results show that both employees and managers can be influenced by own-race bias, and suggest that the own-race bias found in the hiring studies is driven at least partly by discrimination on the part of white job-seekers. In particular, a comparison of the quit rates of employees who receive new managers and those who keep their hiring managers produces evidence that own-race preferences lead white job-seekers to avoid taking jobs with nonwhite managers.[5]

Fourth, we present new evidence on whether own-race bias varies across race groups. Most studies of own-race bias are unable to identify separate biases for different race groups because their analyses are limited to only two race groups (e.g, Dee 2005 and Wolfers and Price 2007). When employing the usual difference-in-difference approach to identifying own-race bias, it is necessary to have more than two race groups in order to identify separate bias parameters for each race group.[6] We know of only two studies that convincingly compare own-race biases across race groups, and these studies reach different conclusions. Giuliano et al. (2009) find evidence of own-race bias in hiring patterns for whites and blacks, and for Hispanics in locations with large Hispanic populations. However, Anbarci and Lee (2008) find no own-race bias among officers who issue speeding tickets; rather, they find that black and Hispanic officers treat whites relatively leniently. Using the same data as Giuliano et al., we too find a general pattern of own-race bias. But because we examine different outcomes than Giuliano et al., we are able to detect anomalies that support Anbarci and Lee's finding that other-race superiors are relatively lenient toward white subordinates.

Finally, we are able to consider the relationship between the effects of own-race bias and the effects of status and identity. We have cases for each outcome where traditional hierarchies are broken, and these norm-breaking relationships allow us to explore the interaction between race and status. We find that in all three cases where our main results show no evidence of own-race bias, our findings could be explained by status and identity effects. Such effects may also help explain why own-race bias appears to be stronger among white employees than employees of other races.

III. Data

A. The Setting

The data are the daily personnel records of a large retail employer from February 1, 1996 through July 31, 1998. These records identify the demographic traits of both managers and their employees at each store, and they give the dates and descriptions of all personnel actions for each individual. We analyze a sample of more than 1,500 store managers who were employed at some point during the 30-month sample period, and more than 100,000 frontline employees who were hired during the sample period.[7]

Our sample contains more than 700 stores located throughout the United States. While geographically diverse, these workplaces nevertheless are all very similar: they are all part of a national chain with highly uniform policies and procedures. In a typical store, there is one full-time manager who has the title “store manager”, and there are 25 to 50 mostly part-time employees.

The managers in our analysis are the “store managers”—the overall manager at each store. The store managers are responsible for all personnel decisions including hiring, dismissals, and promotions. All managers receive a small amount of training in fostering and managing a diverse workforce. Because of frequent managerial turnover and inter-store transfers, 80 percent of the stores have at least one change in management during the 30-month sample period, and roughly 20 percent of all employees get new managers at some point before they leave.[8]

All frontline employees have similar job titles and descriptions. They all rotate through several tasks that involve both dealing with customers and doing support duties. These jobs require only basic skills and employees receive little training. As is common in this sector, employees have very high rates of turnover; the median spell in a store for a frontline employee is 91 days, and roughly 80 percent of employee spells end within a year.

Table 1 summarizes the demographic composition of the managers and employees in our sample. Among managers, 87 percent are white, roughly five percent are black, five percent are Hispanic, and less than three percent are Asian.[9] Managers are also young (the mean age is 30) and predominantly female (78 percent); however, both the age and gender distribution vary across race groups. The average age ranges from 29.2 for Hispanic managers to 31.3 for blacks, and the percent female ranges from 67.3 percent for blacks to nearly 80 percent for whites.

Frontline employees are much more racially diverse than their managers: 64 percent are white, 16 percent are black, ten percent are Hispanic, and seven percent are Asian. Like their managers, these employees are relatively young (mean age 22) and largely female (70 percent). Though the age distribution of employees is similar across race groups, the gender distribution varies substantially—ranging from 56 percent female among Hispanics to 74.5 percent female among whites.

For employees, Table 1 also summarizes firm-specific experience and employment status at the beginning of the employment spell. Roughly 24 percent of all employees are re-hires—that is, they have some prior experience with the firm. However, this statistic varies across race groups, ranging from 18.2 percent for blacks to 25.5 percent for whites. Employment status, on the other hand, is quite similar across race groups. Roughly 65 percent start out with "temporary" status, meaning that they were hired to fill seasonal demand; temporary employees work less than 30 hours per week. Another 33 percent begin their employment spells as part-time (but permanent) employees, and work between 8 and 29 hours per week. Further, about 30 percent of temporary employees become part-time employees, their status routinely being converted to part-time after 80-90 days of employment.[10] Finally, the remaining three percent of employees are classified as full-time and permanent, and they work more than 30 hours per week.

Because our data comes from a single employer, it is important to consider how representative our sample is of a larger population. Our sample is from a retail firm, so perhaps it is most useful to look at how our sample compares to the U.S. retail sector as a whole—a sector that accounts for roughly 18 percent of all U.S. jobs. Compared to the retail sector, our sample is typical with respect to its turnover rates and its racial composition.[11] However, both managers and employees are relatively young (with average ages of 22 and 30 vs. national averages of 32 and 39), and this company has a higher share of both female managers (78 vs. 50 percent) and female employees (70 vs. 66 percent).

Table 2 shows the fraction of employees of each race that work with own-race managers. Whites are most likely to work with own-race managers (90 percent), but the sample contains large numbers of white employees with non-white managers. Among minorities, Hispanics are more likely than blacks (14.9 vs. 8.4 percent) to have own-race managers. The comparison of hiring managers to new managers shows that whites are slightly less likely to have a new white manager than they are to have a white hiring manager, and that blacks and Hispanics are slightly more likely to have a new own-race manager. This pattern reflects an increase in the representation of minorities among managers over the sample period, but is also consistent with the sorting by white job-seekers that is suggested by our analysis.

Our data offers three analytical advantages. First, we can estimate separate own-race effects for each race group because we have many dyads that match employees from each group with managers from each group.[12] We should acknowledge, though, that the estimates for Asians are imprecise because we have fewer cases where Asians are paired with other minorities. Second, we have many stores where new managers are demographically different from those they replace. This within-store variation in manager race allows us to estimate models that control for unobserved differences across stores. Third, because many employees receive new managers, we can test for sorting at the hiring stage by looking at whether own-race bias has a stronger effect under new managers than under hiring managers.

B. Dependent Variables: Quits, Dismissals, and Promotions

The definition of our dependent variables is based on company codes that classify both personnel actions and the reasons for these actions.

1. Exits

Among the frontline employees hired during our 30-month sample period, we observe well over 50,000 exits. Our analysis focuses on two types of exits: job-related quits (54 percent of exits) and dismissals (7 percent of exits). We exclude both market-driven layoffs (9.2 percent of exits) and those who leave voluntarily to move or to return to school (20.3 percent of exits).[13] Quits include voluntary exits that occur because an employee is dissatisfied or has found a better job; those who quit without giving a reason; and those who simply stop showing up for work. Dismissals are involuntary exits that result from tardiness, absenteeism, substandard performance, violations of specific company policies, or various forms of dishonesty.

2. Promotions

The variable we use to analyze promotions is the number of days after hire until the first time an employee is promoted to a new job title. To maintain a sufficient sample size, we pool the 15 different job titles to which an employee may be promoted.[14] In all, we observe roughly 2,500 first-time promotions.

Figure 1 shows Kaplan-Meier estimates of the "failure functions" for quits, dismissals, and promotions. These graphs plot the fraction of employees that have quit, been fired, or been promoted as a function of the number of days since being hired.[15] Separate estimates are constructed for each race group. The first two graphs show that in general, quit and dismissal rates both rise quickly in the first 100 days, and a majority of quits and dismissals occur within a year of employment. In contrast, the graph for promotions shows that very few promotions occur within the first 100 days, and less than five percent of employees receive a promotion within the first year. However, the fraction of employees promoted increases steadily over time.

While the shape of each failure function is similar across race groups, there are clear racial differences in quit, dismissal, and promotion rates at each point in time. Quit rates are largest among blacks and lowest among Asians. Dismissal rates are also highest among blacks, but they are lowest among whites. And promotion rates are highest among whites and Hispanics, and lowest among Asians and blacks. As we'll see below, racial differences in these outcomes persist even after controlling for the employee's location of employment, month of hire, and other employee characteristics.

IV. Empirical Strategy

To examine whether own-race bias affects rates of employee quits, dismissals, and promotions, we estimate continuous time Cox proportional hazard models (Cox 1975) in an independent competing risks framework. Because this approach assumes that the risks are independent conditional on the covariates in the model, we employ several strategies to control for other determinants of our employment outcomes.[16]

First, we eliminate virtually all variation in job characteristics by restricting the analyzed sample to employees with a single job title at a company that has highly uniform workplaces.

Second, we control for observed characteristics of both the employee and the employee's manager. Employee characteristics include demographics (race, ethnicity, age, and gender), and indicators for previous experience with the company and for part-time and temporary (vs. fulltime/permanent) status. Manager characteristics include demographics and an indicator for whether the manager is new (as opposed to the manager who hired the employee).

Third, we control for time-series variation in labor market conditions by including a vector of dummy variables indicating the month in which the employee was hired.

Finally, to control for all fixed characteristics of both the local labor market and the workplace, we estimate stratified models in which each store has its own flexible, baseline hazard function. This approach not only avoids imposing assumptions about the form of the baseline hazard function, it also allows us to control for all fixed characteristics of each store without estimating hundreds of additional parameters.[17] We are able to identify the stratified model because manager turnover in our sample creates ample within-store variation in manager race.

Under the assumption of conditional independence, we can estimate the competing risks model by estimating a separate model for each risk and by treating exits due to the other risks as censored. We begin by estimating a baseline model for each employment outcome that shows whether the likelihood of the outcome is correlated either with employee characteristics or with manager characteristics. For each outcome, the hazard function for employee i in store j is specified as:

(1) $h_{ij}(t) = h_j(t) * exp(X_{ij}\,\beta_X + M_{ijt}\,\beta_M + T_i\,\beta_T)$.

Here $h_j(t)$ is the baseline hazard for store j, and t is the number of days that individual i has been employed at store j. The regressors in this model include variables characterizing the employee (X_{ij}), variables characterizing the employee's current manager (M_{ijt}), and a vector of thirty dummies indicating the sample month in which the employee was hired (T_i).

In equation (1), white is the baseline (omitted) race group. In the dismissals regression, for example, a positive coefficient on the dummy variable indicating that the employee is black would imply that, on average, black employees are more likely to be dismissed than white employees who work in the same store and have similar observables. Such differences in average dismissal rates between black and white employees could be due to some form of bias or to unobserved, racially correlated skills or performance.

The central question is whether own-race bias affects our employment outcomes. We use a difference-in-difference approach and estimate a separate own-race effect for each race group. We add to equation (1) the interaction of each employee race indicator with the manager race indicator corresponding to the same race. This results in four new dummy variables—one for each race group—that indicate whether the manager's race matches the race of the employee. Our estimation equation becomes:

(2) $h_{ij}(t) = h_j(t) \cdot exp(X_{ij}\,\beta_X + M_{ijt}\,\beta_M + T_i\,\beta_T + White_{ij}{\times}MgrWhite_{ijt}{\cdot}\beta_{WW} + Black_{ij}{\times}MgrBlack_{ijt}{\cdot}\beta_{BB}$

$+ Hispanic_{ij}{\times}MgrHispanic_{ijt}\cdot\beta_{HH} + Asian_{ij}{\times}MgrAsian_{ijt}{\cdot}\beta_{AA})$.

The coefficients on the four new own-race dummy variables (β_{WW}, β_{BB}, β_{HH}, and β_{AA}) are the central parameters of interest—they are the estimated effects of having an own-race manager for each race group.

It is important to note that identification of a separate own-race effect for each race group requires having at least three race groups in the analysis. This is because each coefficient β_{kk} is estimated by comparing differences in average outcomes between race k and race j employees under race k and race l managers, where $l \neq j$. The condition that $l \neq j$ is necessary in order to hold constant the own-race dummy variables for race groups j and l, and thereby to isolate the own-race effect for race k. In other words, it is necessary to have managers from an independent race group l to provide a "no-bias" baseline for the differential in employment outcomes between race k and race j employees. Separate own-race effects for race k and race j can then be identified by comparing this baseline differential to the differentials under race k managers and race j managers, respectively.[18]

In analyses with only two race groups, it is not possible to identify a separate own-race bias effect for each group because there is no independent comparison group that can be used to provide a no-bias benchmark. Instead, each of the two groups must be used as a benchmark for the other. As a result, only one "difference-in-difference" can be estimated (that is, only one own-race interaction term can be included in the regression), and the resulting estimate conflates the effects of own-race bias for the two groups being studied. The presence of four groups in our analysis therefore provides an advantage over previous studies that focus on only two groups.

The fact that each coefficient β_{kk} is identified in our study through comparisons involving two independent comparison groups (j and l, where $j \neq l$) should be kept in mind when interpreting our estimates. For example, the coefficient β_{BB} (the effect of own-race bias for blacks) from the dismissals regression will be appropriate for comparing the black-white differential in dismissal rates under black and Hispanic managers. But if one wants to compare the black-white differential in dismissal rates under black and white managers, this will require summing the own-race effects for blacks (β_{BB}) and for whites (β_{WW}).

V. Results

Table 3 reports the results from our analysis of quits, dismissals, and promotions. For ease of interpretation, we report hazard ratios (exponentiated coefficients) instead of the coefficients themselves. For example, a hazard ratio of 1.10 for a dummy variable would imply that the daily rate of quits, dismissals, or promotions is ten percent higher for the indicated group than for the omitted group, and a hazard ratio of 0.90 would imply a rate that is ten percent lower.[19]

Columns 1a, 2a, and 3a contain the results from the baseline model (equation 1). These results show that employee outcomes are highly correlated with employee race, but are less consistently related to manager race. Compared to the quit rate of white employees, quit rates are six percent higher for blacks and 16 percent lower for Asians. Compared to whites, dismissals rates are 2.3 times higher for blacks and 43 percent higher for Hispanics. And compared to whites, promotion rates are 52 percent lower for blacks, 27 percent lower for Hispanics, and 38 percent lower for Asians. While these baseline estimates are not the focus of our study, it is striking that career paths diverge across race groups so early in these low-skill, entry-level positions.

In contrast, there are only two cases where manager race is significantly related to employee outcomes. Under Hispanic managers, the average dismissal rate is 15 percent higher than under white managers. And under black managers, the average promotion rate is forty percent higher than under white managers. As we will see, black managers promote employees at a higher rate on average mainly because they promote black employees at a higher rate.

Our main results are found in columns 1b, 2b, and 3b of Table 3. These show the estimated effects of having an own-race manager on our three employment outcomes.

A. Quits

In quits, we find evidence of modest own-race biases for white, Hispanic, and Asian employees, but not for black employees. The estimates suggest that the relative quit rates of white, Hispanic, and Asian employees are roughly six percent lower under own-race managers. While only the white interaction coefficient is statistically significant, nevertheless we find that the interaction coefficients for whites, Hispanics, and Asians (β_{WW}, β_{HH}, and β_{AA}): (1) are similar in magnitude (with corresponding hazard ratios of roughly 0.94); (2) do not differ significantly from each other (p>.99); and (3) are jointly significant (p<.01).

To help assess the impact of these biases, we first use our estimates to calculate the predicted probability of an employee quitting within one year of being hired for each combination of manager race and employee race (Table 4).[20] Using these figures, we then calculate the average annual quit rate under managers of each race. In a store with a white manager and an average workforce (which is 64 percent white), our estimates suggest the average annual quit rate would be 61.6 percent. If the manager in this store were black instead of white, the average annual quit rate would rise by about three percent (from 61.6 percent to 63.4 percent). This increase is due to the higher quit rate of white employees (remember the black quit rate is unaffected by manager race differences). If the manager were instead Hispanic, the average annual quit rate would rise by only two percent (from 61.6 percent to 62.9 percent). Here the increase in turnover due to the increase in white quit rates is partly offset by a reduction in Hispanic quit rates.[21]

We note that for one group—white employees—these estimates tell only part of the story.

In section 5.1 below, we present evidence that the own-race effect on white quit rates may understate the level of own-race bias among white workers. We find that own-race bias for white employees is substantially larger among those who have received new managers compared to those who still have their hiring managers. This suggests the own-race effect on white quit rates is modest at least partly because white employees who dislike working for non-white managers tend to avoid taking jobs with non-white managers in the first place.

B. Dismissals

In dismissals (Table 3, column 2b), we find a significant own-race bias for black employees and a weakly significant bias for Hispanic employees. The relative dismissal rate for blacks is 19 percent lower under own-race managers, and for Hispanics it is 16 percent lower. A similar effect is found for Asians, but it is not statistically significant. On the other hand, we find that white employees experience a small reverse bias (though it is not significant). The estimate suggests the relative dismissal rate for whites is five percent *higher* under white managers than under other-race managers. Additional tests show that the black, Hispanic, and Asian interaction coefficients do not differ significantly from each other (p=.96); that they are jointly significant (p=.035); and that their mean differs significantly from the white interaction coefficient (p=.04).

To interpret these estimates, we note first that because we find favorable own-race biases for blacks, Hispanics, and Asians, but a zero or reverse bias for whites, this means that the largest differences in dismissal patterns are seen when comparing the relative dismissal rates of minority groups under different minority-race managers. If, for example, we compare the relative dismissal rate of blacks to Hispanics under black and Hispanic managers, these relative rates show the combined effect of both the own-race bias for blacks and the own-race bias for Hispanics. The combined effect is the sum of the black and Hispanic interaction coefficients

($\beta_{BB} + \beta_{HH}$), or the product of the corresponding hazard ratios (which is 0.685). Hence the dismissal rate of blacks relative to Hispanics is 31.5 percent lower under black managers than it is under Hispanic managers.

Alternately, the combined effects of the various own-race biases on dismissal patterns can be seen by examining differences in the predicted probabilities of being dismissed within one year. In Table 4, we see, for example, that black employees are dismissed at substantially higher rates by all types of managers (including blacks), but that the black-nonblack differentials are substantially smaller under black managers. Under white managers, blacks are 2.2 times more likely than whites to be fired within a year (17.3 percent vs. 7.9 percent), but under black managers, blacks are 80 percent more likely than whites to be fired (16 percent vs. 8.7 percent). And under Hispanic managers, blacks are 80 percent more likely than Hispanics to be fired (20.8 percent vs. 11.5 percent), but under black managers, blacks are only 28 percent more likely than Hispanics to be fired (16 percent vs. 12.5 percent).

C. Promotions

In promotions (Table 3, column 3b), we find the same pattern as in dismissals—that is, an own-race bias for blacks and Hispanics, but a reverse bias for whites. Evidence of own-race bias is again strongest for blacks. The coefficient on the own-race interaction term for blacks is positive and significant, and suggests the relative promotion rate of blacks is 79 percent higher under black managers. For Hispanics, the estimate suggests their relative promotion rate is 40 percent higher under Hispanic managers. Although the coefficient on the own-race interaction term for Hispanics is smaller and not statistically significant, a Wald test cannot reject equality of the coefficients for blacks and Hispanics (p=.50), and together, they are jointly significant (p=.08). Finally, the interaction coefficient for whites, while not significant, does differ significantly from the mean of the black and Hispanic coefficients (p=.06) and indicates a reverse bias. The estimate suggests that whites are promoted at a 20 percent *lower* rate under white managers.

Table 4 again shows the predicted probabilities of being promoted within one year. As in dismissals, blacks have the worst outcomes as a group under all types of managers, but the black-nonblack differentials are substantially smaller under black managers. Under white managers, for example, whites are 4.7 times more likely than blacks to be promoted within a year (4.7 percent vs. 1.0 percent); but under black managers they are 3.3 times more likely (6.2 percent vs. 1.9 percent). Again we see even more dramatic differences when comparing black-Hispanic differentials under black and Hispanic managers. Under Hispanic managers, Hispanics are 1.9 times more likely than blacks to be promoted (2.3 vs. 0.8), but under blacks they are only ten percent more likely (2.1 percent vs. 1.9 percent).

In sum, we find a consistent pattern of own-race bias across all three outcomes and across both employees and managers. But there are anomalies in this pattern. In quits (employee decisions), we find evidence of modest own-race biases among white, Hispanic, and Asian employees. But we do not find bias among black employees. In dismissals and promotions (managerial decisions), we find a common pattern. On one hand, we find evidence of managerial bias in that black, Hispanic, and Asian employees are relatively less likely to be dismissed and more likely to be promoted when their manager is the same race. However, we also find that white employees have similar or *better* relative outcomes under minority managers than under white managers.[22]

VI. Robustness Tests and Extensions

A. Hiring Managers vs. New Managers: A Test for Selection at the Hiring Stage

If managers and employees indulge their racial preferences when they select whom to hire or where to take a job, discrimination is less likely to be found in post-hire outcomes. We are particularly concerned about whether employee selection at the hiring stage reduces the magnitude of the own-race effects on quit rates. Whereas managers face legal and social constraints against selecting employees based on racial preferences, job-seekers face no legal obligation to accept jobs with other-race managers. Hence employee selection might be especially prevalent.

To test for employee selection, we compare the effects of own-race bias on quit rates for two groups—employees who still have their hiring managers and employees with new managers. The test is based on the strong contrast between the potential impact of selection in the two groups. Employees who have their hiring managers were able to select their managers with few constraints, and so selection might minimize the effect of own-race bias on quits in this group. In contrast, employees with new managers did not select their manager, and so our estimates for this group should be free of selection bias. Hence if employee selection is affecting our estimates of own-race bias on quits, we might expect the estimates to be larger among employees who receive new managers.

We find evidence of selection bias for white employees, but not for employees of other races (Table 5, columns 1-2). For white employees, the effect of having an own-race manager on their quit rate is significantly larger (p=.06) when the manager in question is new. Among white employees who still have their hiring managers, the relative quit rate of whites is only four percent lower under white managers compared to non-white managers. Among white employees who have received new managers, there is a fourteen percent difference. We interpret this result as evidence that white employees who dislike working for non-white managers often avoid working for such managers in the first place, and that when such whites involuntarily find themselves working for a non-white manager, their quit rate increases substantially.

To better understand the potential impact of this bias, we calculate the implied effect that a new manager's race has on average annual quit rates in a store that hires a new manager. Because the new manager's race matters most for white employees, we consider the extreme case of a store with all white employees. In this case, our estimates imply that 66.3 percent of employees would quit within a year of receiving a new white manager, while 71.8 percent would quit if the new manager were not white.[23]

Managerial selection can occur either when managers decide whom to hire or when they decide where to take a job. Either type of selection could reduce the effect of own-race bias on dismissals and promotions. Unfortunately, the method used to test for employee selection does not provide a clear test for managerial selection. For employees, our test was based on the contrast between the potential effects of selection among employees with hiring managers (free selection) and employees with new managers (no selection). For managers, the same contrast does not exist between hiring managers and new managers. Among hiring managers, selection may be constrained by legal and social curbs on hiring discrimination; and among new managers, selection may occur because they have some choice over whom to work with when they accept a job and become "new managers."

Still, we perform the analysis of dismissals and promotions to see if any contrasts emerge. Table 5 reports the results. In promotions, we find no significant differences between hiring managers and new managers, but the standard errors are quite large. In dismissals, however, there are two cases where there is a divergence between the two groups of managers. First, the own-race bias effect for blacks is found only for hiring managers; it disappears for new managers. Second, while among hiring managers we find no own-race bias effect for whites, there appears among new managers a significant reverse bias for whites. One interpretation of these differences is that selection bias due to managerial selection on own-race preferences is greater when managers are new.[24] In section 6, we discuss alternate interpretations of these results.

B. Own-race Managers vs. Own-race Coworkers

Another concern is the possibility that employment outcomes depend not only on whether an employee is the same race as the manager, but also on the number of coworkers who are the same race.[25] Because the race of managers and their employees tend to be correlated, our estimates might conflate these two effects. For example, in the cases where we find higher quit rates when the manager is a different race, the estimates could be driven partly by employees who dislike the concomitant racial composition of their coworkers. Similarly, our results regarding dismissals and promotions could reflect the effect that being different from one's coworkers may have on an employee's productivity. To address these possibilities, we re-estimate our hazard models with additional controls for the fraction of an employee's coworkers who are the same race. The results are presented in Table 6.

First, the analysis of quits shows that having more coworkers of the same race does significantly reduce quit rates, especially for white and Asian employees. Moreover, we find that controlling for coworker similarity does somewhat reduce the estimated effects of manager-employee similarity, and that the coefficient on the own-race interaction for whites is no longer statistically significant (Table 6, column 1). However, when whites, Hispanics, and Asians are pooled, the estimated effect of having an own-race manager remains statistically significant and suggests about a four percent reduction in relative quit rates for these three groups (vs. a six percent reduction in our original estimates). Second, the analysis of dismissals and promotions shows no strong relationship between these outcomes and the fraction of employees who share the employee's race. Hence we conclude that our original estimates for these two outcomes reflect the effects of being similar to one's manager, and not the effects of being similar to one's coworkers.

C. Reasons for Quitting—Dissatisfaction with Supervisor

Our firm established a list of reasons for exiting and recorded the main reason for each exit. One reason for quitting is "dissatisfaction with supervisor." Here we look at whether having an own-race manager affects the likelihood that an employee cites this as the main reason for quitting.[26] If having an own-race manager reduces quit rates because of racial bias in the way employees view their managers, we might expect that having an own-race manager will also reduce the likelihood that employees list "dissatisfaction with supervisor" as their reason for quitting.

We test this prediction by estimating a linear probability model for the sample of employees who quit with regressors similar to those used in the hazard models above. We find a significant own-race bias only for white employees (Table 7). While roughly 0.6 percent of white employees with white managers report being dissatisfied with their supervisors, this figure is roughly double for whites with non-white managers. Further analysis (results not shown in table) reveals that the coefficient on the own-race interaction for whites is larger for those with new managers (-0.010 vs. -0.006); however the difference between those with new managers and those with hiring managers is not statistically significant in this case.

VII. Discussion

Our main results show a consistent pattern of own-race bias across all outcomes and across both employees and managers. Hence we find strong support for in-group theories. But the anomalies in our results suggest that when considering the effects of racial bias in the workplace, one must also take into account the effects entailed by status and identity theories. Here we discuss how our findings support these conclusions, and consider the specific behaviors our results may reflect.

In quits, we find modest effects of own-race biases among white, Hispanic, and Asian employees. These biases could reflect employee preferences either for own-race managers or for workplace environments provided by same-race managers. Or they could reflect the response of employees to more favorable treatment by own-race managers. The one anomaly is that we find no own-race effect for blacks. Why might the race of the manager have no effect on the likelihood that blacks will quit? Status and identity theories suggest that because blacks have historically held a lower position in U.S. society and have worked mainly for non-black managers, it is possible they have become more accustomed to working for other-race managers.

Further analysis of quits suggests that own-race bias is strongest among white employees. First, for whites alone, having a different-race manager doubles the likelihood that "dissatisfaction with supervisor" is reported as the reason for quitting. Moreover, the own-race effect on quit rates for whites is especially large among those with new managers, and this suggests that a preference for own-race managers also determines where white job-seekers accept jobs in the first place. Why would own-race bias be stronger among whites than other race groups? Our finding that whites are, if anything, treated more favorably by non-white managers suggests that such a bias is not a response to differential treatment by managers. However, identity theory suggests a stronger bias could come from a combination of two things.

Whites may be subject both to own-race bias and to a status threat produced by a non-traditional hierarchy (minorities managing whites).

Dismissals and promotions are managerial decisions, and in both outcomes we find the same pattern of results. Blacks, Hispanics, and Asians have better relative outcomes under own-race managers, but just the opposite is true of white workers. Why do minority managers dismiss own-race employees less often and promote them more often? Such own-race biases may reflect discriminatory preferences or efficiency motives. These explanations are not mutually exclusive, and we cannot distinguish between them. In fact, it may be impossible to distinguish them empirically because discriminatory preferences may be an underlying cause of productivity differences (for example, see models by Charles 2000 and McLeod 2003).

If the biases are driven by discriminatory preferences, which race groups are doing the discriminating? There are four possible patterns of managerial discrimination consistent with our results: (1) Black (Hispanic, Asian) managers could favor black (Hispanic, Asian) employees; (2) Black (Hispanic Asian) managers could penalize non-black (non-Hispanic, non-Asian) employees; (3) non-black (non-Hispanic, non-Asian) managers could favor non-black (non-Hispanic, non-Asian) employees; (4) non-black (non-Hispanic, non-Asian) managers could penalize black (Hispanic, Asian) employees.

Again, these behaviors are not mutually exclusive, and we cannot distinguish between them. To be sure, previous literature on employment discrimination has found evidence of discrimination against blacks especially, and against Hispanics to a lesser degree.[27] This is consistent with our results in that own-race bias is strongest for blacks, and it is plausible that our results are driven at least partly by discrimination against blacks (Hispanics) on the part of non-black (non-Hispanic) managers.

In dismissals and promotions, our second basic finding is that white employees with other-race managers have relative outcomes that are similar or *better* than the outcomes of whites with white managers. This result cannot be explained by theories of own-race bias. However, status and identity theories provide a plausible explanation. These theories suggest non-white managers may find that white employees are less willing to accept their authority and are more likely to quit. What's more, non-white managers who exercise authority over whites may find that such norm-breaking relationships cause anxiety and psychological discomfort. Hence even if non-white managers hold own-race biases, these factors might still lead them to be relatively deferential toward whites and so dismiss whites less often and promote them more often.

If we accept such status effects, it is worth noting that an own-race coefficient of zero for whites does not necessarily mean there is no own-race bias among whites. Rather, it could reflect a combination of effects that are equal in magnitude and opposite in sign—that is, a combination of white managers favoring whites (an own-race bias) and non-white managers favoring whites (a status/identity effect or reverse bias).

In dismissals and promotions, we also compared the effects of own-race bias for hiring managers and new managers, and in dismissals we found two cases where there was a divergence between these two groups. First, while there was a strong own-race bias for blacks among hiring managers, this bias disappears among new managers. Second, while among hiring managers we find no own-race bias for whites, there appears among new managers a significant reverse bias for whites. Why would the own-race bias for blacks disappear among new managers? And why would a reverse bias for whites appear?

We propose that status, again, may play an important role. We argued above that a manager's ability to command authority may depend on the social status associated with his race group relative to the race group of the employees he supervises. Here another type of status difference may also be important—the difference in status between managers supervising employees whom they hired *vs*. new managers supervising incumbent employees. If new managers are less secure in their authority over incumbent employees, they may be less likely to indulge racial biases in decisions about these employees. Such a status effect could explain the fact that the own-race bias for blacks disappears among new managers.

A status disadvantage for new managers could also help produce a *reverse* bias for whites. It could do so either through its effect on new white managers, through its effect on new, non-white managers, or through a combination of both. First, consider the case of new white mangers. We noted above that the own-race coefficient of zero for white employees with hiring managers could reflect a combination of two offsetting effects—white managers favoring whites (an own-race bias) and non-white managers favoring whites (a reverse bias/status effect). If so, then a reluctance by new white managers to indulge own-race bias could cause the status effect to prevail among new managers.

The case of new, non-white managers is more complex. We recall from our analysis of quits that incumbent white employees are especially likely to quit when they receive new, non-white managers. These managers thus face a particular need to retain white employees. What's more, new non-white managers have two status disadvantages in relation to white employees. Not only are they minorities managing whites, they also have the lesser status of new managers. Hence new non-white managers must supervise a group of white employees who are especially likely to quit on them, and they must do so while they are at a double status disadvantage. Such a situation could make new non-white managers especially deferential toward whites.

VIII. Conclusion

We examine how racial matches between managers and employees affect rates of employee quits, dismissals, and promotions. While the racial diversity of our company reflects the changing nature of the American workforce, our findings point to the enduring nature of both racial preferences and traditional hierarchies. First, we find a general pattern of own-race bias across all our outcomes and across both employees and managers. But second, we also find variations in this pattern, and these variations suggest that when considering the effects of race, one must also consider the effects of status and identity.

Quits are employee decisions, and our quit results produce evidence of employee bias. We find modest own-race biases for white, Hispanic, and Asian employees, though not for black employees. However, further analysis of quits suggests that own-race bias is strongest among whites. We find that unless white employees choose their managers, they are much more likely than other employees to quit under different-race managers. This suggests that white job-seekers sort into workplaces based on preferences for own-race managers, and hence that the own-race effect on quit rates may understate the level of own-race bias among white workers. A stronger own-race bias among white workers could occur if they feel their status threatened when they work for minority managers.

Dismissals and promotions are managerial decisions, and in both we find a similar pattern. On one hand, we find evidence of managerial own-race bias in that black, Hispanic, and Asian employees are relatively less likely to be dismissed and more likely to be promoted when their manager is the same race. But remarkably, we also find that white employees have similar or *better* relative outcomes under minority managers than under white managers. Such deferential treatment of whites by minority managers could be explained by the status disadvantages minority managers face in non-traditional hierarchies.

References

Akerlof, George and Rachel Kranton. 2000. “Identity and Economics.” *Quarterly Journal of Economics* 115(3): 715-53.

Anbarci, Nejat and Jungmin Lee. 2008. "Speed Discounting and Racial Disparities: Evidence from Speeding Tickets in Boston." *IZA Discussion Papers* 3903, Institute for the Study of Labor (IZA).

Antonovics, Kate and Brian G. Knight. 2009. “A New Look at Racial Profiling: Evidence from the Boston Police Department.” *Review of Economics and Statistics* 91(1): 163-77.

Athey, Susan; Christopher Avery; and Peter Zemsky. 2000. “Mentoring and Diversity.” *American Economic Review* 90(4): 765-786.

Arrow, Kenneth. 1973. “The Theory of Discrimination.” In *Discrimination in Labor Markets*, eds., O. Ashenfelter and A. Rees, 3-33. Princeton, NJ: Princeton University Press.

Bates, Timothy. 1994. “Utilization of Minority Employees in Small Business: A Comparison of Nonminority and Black-Owned Enterprises.” *Review of Black Political Economy* 23(1): 113-121.

Becker, Gary S. 1971. *The Economics of Discrimination, 2nd edition*. Chicago, IL.: The University of Chicago Press.

Carrington and Troske. 1998. “Interfirm Segregation and the Black/White Wage Gap.” *Journal of Labor Economics* 16(2): 231-60.

Charles, Kerwin Kofi. 2000. “A Simple Model of Subtle Discrimination.” University of Michigan. Unpublished.

Cox, David R. 1975. “Partial Likelihood.” *Biometrika* 62(2): 269-276.

Dee, Thomas S. 2005. "A Teacher Like Me: Does Race, Ethnicity or Gender Matter?" *American Economic Review*. 95(2): 158-165.

DeVaro, Jed; Ghosh, Suman; and Cindy Zoghi. 2007. "Job Characteristics and Labor Market Discrimination in Promotions: New Theory and Empirical Evidence." *Cornell University: ILR Collection Working Papers*.

Dononue, John and Steven Levitt. 2001. "The Impact of Race on Policing and Arrests." *Journal of Law and Economics* 44(2): 367-394.

Epstein, Richard. A. 1992. *Forbidden Grounds: The Case against Employment Discrimination Laws*. Cambridge, MA, Harvard University Press.

Giuliano, Laura, David I. Levine, and Jonathan Leonard. 2009. "Manager Race and the Race of New Hires." *Journal of Labor Economics* 27(4): 589-631.

Lang, Kevin. 1986. "A Language Theory of Discrimination." *Quarterly Journal of Economics* 101(2): 363-82.

Leonard, Jonathan and David I. Levine. 2006. "The Effects of Diversity on Turnover: A Very Large Case Study." *Industrial and Labor Relations Review* 39(4): 457-572.

MacLeod, Bentley. 2003. "Optimal Contracting with Subjective Evaluation." *American Economic Review* 93(1): 216-240.

Price, Joseph and Justin Wolfers. 2007. "Racial Discrimination among NBA Referees." Working Paper No. 13206. Cambridge, MA: National Bureau of Economic Research.

Ridgeway, Cecilia. 2007. "Status Construction Theory." In *The Blackwell Encyclopedia of Sociology* Vol. IX, ed., G. Ritzer, 4756-4759. Oxford: Blackwell.

Sørensen, Jesper B. 2004. "The Organizational Demography of Racial Employment Segregation." *American Journal of Sociology* 110(3): 626-71.

Stauffer, Joseph M. and M. Ronal Buckley. 2005. "The Existence and Nature of Racial Bias in Supervisory Ratings." *Journal of Applied Psychology* 90(3): 586-591.

Stoll, Michael A., Steven Raphael, and Harry J. Holzer. 2004. "Black Job Applicants and the Hiring Officer's Race." *Industrial and Labor Relations Review* 57(2): 267-287.

Thomas, David A. 1990. "The Impact of Race on Managers' Experiences of Developmental Relationships." *Journal of Organizational Behavior* 2(4): 479-492.

Wesolowski, Mark A. and Kevin W. Mossholder. 1997. "Relational Demography in Supervisor-Subordinate Dyads: Impact on Subordinate Job Satisfaction, Burnout, and Perceived Procedural Justice." *Journal of Organizational Behavior* 18(4): 351-362.

Notes

1. Specifically, we estimate stratified hazard models in which each store has its own baseline hazard function. The models are identified through within-store variation in manager race that results when new managers are a different race from those they replace.

2. Two models should be noted that combine taste-based and efficiency-based factors by positing that managerial bias or the perception of such bias *can cause* differences in productivity. Charles (2000) argues biased managers may sabotage employees, and MacLeod (2003) suggests employees may reduce their effort if they think that managers are biased in their evaluations.

3. Charles (2000) and MacLeod (2003) suggest that under biased managers, other-race employees may also be less productive because of discriminatory sabotage or perceived biases in evaluation. Footnote 2 explains these ideas.

4. Carrington and Troske (1998) and Bates (1994) also find that black managers and business-owners employ more blacks; however these studies are less able to distinguish own-race bias from other sources of segregation.

5. See section 5.1 for details of this analysis and the results.

6. See section 3 below for a full explanation of this point. It is worth noting that Dee (2005) and Price and Wolfers (2007) both present separate estimates of own-race bias by race group despite having only two race groups. However, while both papers employ the standard difference-in-difference approach in their main analysis, this method is dropped when estimating the separate own-race biases for each race group. As a result, the biases are identified only under the questionable assumptions that in the absence of bias, evaluations of students would be uncorrelated with teacher race (in Dee 2005) or that foul-calling rates would be uncorrelated with referee race (in Price and Wolfers 2007).

7. For employees, we exclude left-censored employment spells (those who were hired before Feb. 1, 1996) because we lack dates of hire for these employees. When analyzing promotions, we restrict the sample to those with no prior company experience (that is, we exclude re-hires).

8. Managerial spells at a store are somewhat more likely to end in transfers to other stores rather than in exits from the employer.

9. The remaining employees are classified as Native American or "other". The company's records classify Hispanics by ethnicity and not by race; hence these categories are mutually exclusive and collectively exhaustive.

10. Conditional on tenure, there are no significant racial differences in the rate of transition from temporary to part-time status.

11. The turnover comparison is based on estimates from the National Longitudinal Survey of Youth 1997. Among those 16-20 year olds who worked in low-wage (≤$9.00/hr) retail jobs in 1999, the median employment spell was about 110 days, and 87 percent left their job within a year. The racial composition comparison is based on based on all individuals in the 1996-1998 monthly Current Population Survey who had retail jobs. Of retail managers, 81 percent were white, 7 percent black, 7 percent Hispanic, and 5 percent Asian; retail employees were 73 percent white, 13 percent black, 10 percent Hispanic, and 4 percent Asian.

12. For example, we observe more than 1,000 Hispanic employees with black managers and more than 2,000 black employees with Hispanic managers.

13. Analyses of layoffs and school-related exits showed no evidence of own-race bias in these outcomes. (Results are available from the authors on request.) As few stores close in our sample period, layoffs are typically due to the end of the holiday shopping season. The remaining 9.5 percent of spell terminations are due either to within-company transfers or to leaves of absence,

and so are not separations from the employer. We do not analyze transfers because we lack information on the reason for the transfer. For example, we cannot distinguish among transfers requested by the employee due to friction with the manager, those that were tantamount to promotions (such as relocations to a more desirable location), and those that resulted simply from a change in the employee's place of residence.

14. The lowest paid of these 15 jobs (which accounts for roughly one third of the promotions) earns on average 12 percent more than the entry level job. The highest paid of these 15 jobs earns 26 percent more than the lowest. We find no evidence that either the *type* of job code at first promotion or the increase in pay is affected by manager-employee similarity. However, our sample is not large enough to allow precise estimates of these relationships.

15. In each case, exits due to competing risks are treated as censored. For example, the fraction that quit at time *t* represents the quit rate in the absence of other exit risks.

16. Our method eliminates unobserved heterogeneity across time and across store locations, but not across individuals within stores. In hazard models, even if the omitted variables are not correlated with the covariates of interest, unobserved individual-level heterogeneity is a concern because of the bias that results from higher-risk individuals being selected out of the sample more quickly. Also, such selection bias may be exacerbated if the omitted variables cause residual correlation across competing risks. We therefore investigated the importance of unobserved individual heterogeneity for our results by estimating "frailty" models in which the survival time of each individual depends on a random effect and the baseline hazard is assumed to have a Weibull distribution. The results (available from the authors) were qualitatively similar to those presented below, but often implied larger own-race effects. Hence, they suggest that the estimates presented in this paper might be biased toward zero. However, the frailty/random

effects model also has two important limitations. First, it relies more heavily on functional form assumptions; and second, it does not allow for stratification by store location. For these reasons, we focus on the results from the stratified Cox model described below. For similar reasons, we did not attempt to account for residual correlation among the competing risks in our model.

17. Here, the likelihood function is formed by first calculating for each duration time *t* the conditional probability that, of all individuals employed at a given store for at least *t* days, a particular individual *i* exits (or is promoted) on day *t*; and by then taking the product of these conditional probabilities (Cox 1975). We use the Breslow (1974) method for handling ties.

18. Because we have four race groups, each race group *k* has available six combinations of employee race *j* and manager race *l* such that $j \neq l$. Hence each own-race estimate βkk is based on six such difference-in-difference comparisons.

19. The standard errors reported in the tables are computed using the delta rule; that is, they are the standard errors of the coefficients multiplied by the exponentiated coefficients. A test of significance is a test of whether the hazard ratio differs from 1.00.

20. These probabilities are obtained by first estimating the cumulative hazard function for white employees with white managers. The probabilities for the other groups are then obtained by adjusting the hazard function by the appropriate coefficients from Table 3, columns 1b, 2b, and 3b. Because exits due to competing risks are treated as censored, these figures represent net or "cause-deleted" probabilities (for example, the probability of quitting in the absence of other exit risks).

21. Note that the coefficients on the black manager and Hispanic manager indicators in Table 3, column 1b indicate that there is no change in average quit rates associated with having a black or Hispanic manager *other than* the changes associated with own-race bias. Hence our calculated

differences in average quit rates across white, black and Hispanic managers reflect only the effects of own-race bias. For the case of Asians, the relevant coefficients suggest that Asian managers have lower average quit rates even among non-Asian employees and especially low quit rates among Asian employees. However, we do not emphasize the results for Asians because of the relatively large standard errors.

22. These findings confirm the importance of estimating a separate own-race bias for each race group. Indeed, when we perform the analysis using only whites and blacks in our data, the resulting estimates conflate the own-race effects for blacks and whites and therefore obscure our main findings. Specifically, these estimates imply zero own-race bias in quit rates, and moderate but only marginally significant own-race biases in dismissals and promotions.

23. These probabilities are calculated similarly to those in Table 4. We obtain the annual quit rate for white employees with new white managers by estimating the cumulative hazard function for this group of employees. Because having a new manager increases the white quit rate, this probability (66.3 percent) is higher than the figure for white employees with white managers in Table 4 (62.2 percent). Then, the annual quit rate of white employees with new, non-white managers is obtained by adjusting the hazard function by the relevant coefficient from Table 5. As before, exits due to competing risks are treated as censored, so these estimates represent the probabilities of quitting within one year in the absence of other exit risks.

24. This would imply that managers' racial preferences have a greater impact on where they take jobs than on whom they hire once they have a job. For example, managers with a distaste for supervising other-race employees might avoid taking jobs where other-race workers are already employed; but may be compelled to hire other-race employees once they accept a job.

25. Levine and Leonard (2006) and Sorensen (2004) both find that employee turnover rates are inversely related to the level of same-race representation among co-workers.

26. The fraction of employees who indicate this reason is small—only 0.7 percent of all job-related quits. The bulk of employees who quit are vague about their reasons for doing so. Roughly 35 percent give no reason, 33 percent give "other" as the reason, and 20 percent say only that they found a better job.

27. Audit studies by Turner et al. (1991) and Bendick et al. (1994) find that black research assistants posing as job applicants ("testers") receive fewer job offers than white testers with similar qualifications. Kenney and Wissoker (1994) find similar results for Hispanics. Bertrand and Mullainathan (2004) find that resumes with black-sounding names obtain fewer interview offers than identical resumes with white-sounding names. DeVaro, Ghosh, and Zoghi (2007) find that non-whites are promoted at lower rates than similarly qualified whites. None of these studies analyzes the role of manager race or preferences in such discrimination, either because there is no data on manager traits or because there is not enough variation.

Table 1
Manager and Employee Characteristics by Race

	Percent of Total	Percent Female	Age	Percent with Firm Experience	Percent Part-Time	Percent Temporary
Managers						
White	87.0	79.6	30.4 (5.1)			
Black	4.8	67.3	31.3 (5.1)			
Hispanic	5.5	70.9	29.2 (4.1)			
Asian	2.4	75.4	29.4 (4.4)			
Other	0.3	98.3	27.5 (4.0)			
All	100.0	78.4	30.3 (5.1)			
Employees						
White	64.4	74.5	22.1 (7.2)	25.5	32.8	64.3
Black	16.4	64.8	22.7 (6.3)	18.2	32.6	64.9
Hispanic	10.3	55.9	22.1 (5.9)	21.1	32.6	63.6
Asian	6.9	65.3	20.9 (5.0)	25.1	32.5	64.8
Other	1.9	75.2	22.2 (6.2)	23.8	31.1	66.6
All	100	70.4	22.1 (6.8)	23.8	32.7	64.4

Note: Based on sample of N>100,000 employees hired between February 1, 1996 and July 31, 1998, and N>1,500 managers employed during this period. Age, experience and employment status are measured at the beginning of the employment spell. Parentheses contain standard deviations of age.

Table 2
Share of Employees with Own-race Managers, by Employee Race

	Employees with Hiring Managers	Employees with New Managers	All
White employees	90.3	88.2	89.9
Black employees	8.2	9.0	8.4
Hispanic employees	14.6	16.1	14.9
Asian employees	5.3	4.8	5.2
N >	100,000	20,000	120,000

Note: Figures are the percent of employees in each category whose manager is the same race or ethnicity.

Table 3
Estimates from Cox Proportional Hazard Models of Quits, Dismissals, and Promotions

	Quits		Dismissals		Promotions	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Employee is Black	1.063**	1.002	2.290**	2.421**	0.481**	0.375**
	(0.012)	*(0.029)*	*(0.061)*	*(0.166)*	*(0.040)*	*(0.076)*
Employee is Hispanic	0.981	0.933*	1.431**	1.511**	0.732**	0.587**
	(0.013)	*(0.028)*	*(0.048)*	*(0.107)*	*(0.064)*	*(0.114)*
Employee is Asian	0.841**	0.797**	1.072	1.127	0.621**	0.650
	(0.014)	*(0.024)*	*(0.047)*	*(0.085)*	*(0.076)*	*(0.403)*
Current manager is Black	1.043	1.000	1.006	1.128	1.402**	1.073
	(0.029)	*(0.037)*	*(0.070)*	*(0.099)*	*(0.174)*	*(0.215)*
Current manager is Hispanic	1.016	0.997	1.151‡	1.223*	1.055	0.843
	(0.029)	*(0.035)*	*(0.084)*	*(0.100)*	*(0.141)*	*(0.182)*
Current manager is Asian	0.997	0.968	1.136	1.194‡	1.121	0.933
	(0.038)	*(0.042)*	*(0.104)*	*(0.120)*	*(0.328)*	*(0.347)*
White employee*White manager		0.936*		1.047		0.801
		(0.027)		*(0.074)*		*(0.154)*
Black employee*Black manager		1.024		0.812*		1.791*
		(0.043)		*(0.074)*		*(0.503)*
Hispanic employee*Hispanic manager		0.937		0.844‡		1.408
		(0.042)		*(0.086)*		*(0.404)*
Asian employee*Asian manager		0.939		0.820		0.990
		(0.072)		*(0.140)*		*(0.615)*

Note: Hazard ratios from Cox proportional hazard models, stratified by store. Model of promotions is also stratified by twelve age categories. Additional controls are: dummy variables indicating employee gender, manager gender, an employee part-time status, employee age and age squared (at time of hire), manager age, and a dummy variable indicating whether the manager is new (that is, not the manager who hired the employee). Robust standard errors in parentheses, adjusted for clustering on employee. ‡ Significant at 10 percent; * significant at 5 percent; ** significant at 1 percent (based on test that the hazard ratio differs from one). Tests of equality and joint significance on the own-race interaction coefficients indicate the following: *Quits*: The White, Hispanic, Asian interaction coefficients do not differ significantly (p>.99); they are jointly significant (p<.01); and their mean differs significantly from the Black interaction coefficient (p=.07). *Dismissals*: The Black, Hispanic, and Asian interaction coefficients do not differ significantly (p=.96); they are jointly significant (p=.035); and their mean differs significantly from the White interaction coefficient (p=.05). *Promotions*: The Black and Hispanic interaction coefficients do not differ significantly (p=.50); they are jointly significant (p=.08); and their mean differs significantly from the White interaction coefficient (p=06).

Table 4
Predicted Probabilities of Quitting, Being Dismissed, and Being Promoted within One Year of Being Hired.

One-Year Quit Probabilities:

	White Manager	Black Manager	Hispanic Manager	Asian Manager
White employee	0.622	0.646	0.645	0.635
Black employee	0.614	0.623	0.613	0.602
Hispanic employee	0.621	0.621	0.596	0.609
Asian employee	0.562	0.562	0.561	0.528

One-Year Dismissal Probabilities:

	White Manager	Black Manager	Hispanic Manager	Asian Manager
White employee	0.079	0.087	0.092	0.090
Black employee	0.173	0.160	0.208	0.203
Hispanic employee	0.112	0.125	0.115	0.132
Asian employee	0.085	0.095	0.103	0.083

One-Year Promotion Probabilities:

	White Manager	Black Manager	Hispanic Manager	Asian Manager
White employee	0.047	0.062	0.049	0.054
Black employee	0.010	0.019	0.008	0.009
Hispanic employee	0.020	0.021	0.023	0.018
Asian employee	0.031	0.033	0.026	0.029

Note: Probabilities are obtained by first estimating the cumulative hazard function for white employees with white managers, and then adjusting the hazard function by the appropriate coefficients from Table 3, columns 1b, 2b, and 3b to obtain the probabilities for the other groups. Because exits due to competing risks are treated as censored, these figures represent net or “cause-deleted” probabilities (for example, the probability of quitting in the absence of other types of exit risk).

Table 5
Effects of Own-Race Bias for Hiring Managers vs. New Managers

	Quits		Dismissals		Promotions	
	hazard ratio	χ^2 (Pr> χ^2)	hazard ratio	χ^2 (Pr> χ^2)	hazard ratio	χ^2 (Pr> χ^2)
White employee*White manager*Hiring manager	0.957		0.997		0.803	
	(0.030)	3.42‡	*(0.073)*	4.15*	*(0.158)*	0.00
White employee*White manager*New manager	0.858**	*(0.064)*	1.319*	*(0.042)*	0.801	*(0.986)*
	(0.049)		*(0.178)*		*(0.168)*	
Black employee*Black manager*Hiring manager	1.013		0.784*		1.734‡	
	(0.048)	0.20	*(0.075)*	2.16	*(0.563)*	0.04
Black employee*Black manager*New manager	1.057	*(0.653)*	0.988	*(0.141)*	1.892	*(0.840)*
	(0.096)		*(0.152)*		*(0.736)*	
Hispanic employee*Hispanic manager*Hiring manager	0.935		0.857		1.222	
	(0.044)	0.42	*(0.095)*	0.14	*(0.433)*	0.61
Hispanic employee*Hispanic manager*New manager	0.955	*(0.517)*	0.791	*(0.706)*	1.644	*(0.436)*
	(0.051)		*(0.155)*		*(0.552)*	
Asian employee*Asian manager*Hiring manager	0.946		0.817		0.756	
	(0.076)	2.45	*(0.147)*	0.00	*(0.805)*	0.10
Asian employee*Asian manager*New manager	0.890	*(0.118)*	0.836	*(0.956)*	1.116	*(0.748)*
	(0.076)		*(0.337)*		*(0.786)*	

Note: Hazard ratios from Cox proportional hazard models with control variables as in Table 3, plus all interactions of manager race and employee race indicators with an indicator that the manager is new. Robust standard errors in parentheses, adjusted for clustering on employee. ‡ Significant at 10 percent; * significant at 5 percent; ** significant at 1 percent (based on test that the hazard ratio is different from one). Second column for each set of results reports Wald test of equality for each pair of hazard ratios.

Table 6
Own-Race Managers vs. Own-Race Coworkers

	Quits	Dismissals	Promotions
White * Percent white coworkers	0.746**	1.015	1.124
	(0.029)	*(0.113)*	*(0.255)*
Black * Percent black coworkers	0.899‡	0.974	1.162
	(0.055)	*(0.140)*	*(0.516)*
Hispanic * Percent Hispanic coworkers	0.971	0.937	0.443‡
	(0.073)	*(0.169)*	*(0.191)*
Asian * Percent Asian coworkers	0.727**	0.755	0.970
	(0.089)	*(0.251)*	*(0.746)*
Employee white, manager white	0.964	1.045	0.675‡
	(0.028)	*(0.075)*	*(0.137)*
Employee black, manager black	1.026	0.815*	2.165**
	(0.044)	*(0.075)*	*(0.631)*
Employee Hispanic, manager Hispanic	0.945	0.852	1.576
	(0.043)	*(0.089)*	*(0.478)*
Employee Asian, manager Asian	0.983	0.857	1.052
	(0.076)	*(0.152)*	*(0.677)*

Note: Hazard ratios from Cox proportional hazard models with control variables as noted in Table 3. Robust standard errors in parentheses, adjusted for clustering on employee. ‡ Significant at 10 percent; * significant at 5 percent; ** significant at 1 percent.

Table 7
Probability that "Dissatisfaction with Supervisor" is the Employee's Reason for Quitting

	(1)	(2)
Constant	0.004	0.011*
	(0.004)	*(0.005)*
Employee is Black	0.001	-0.006*
	(0.001)	*(0.002)*
Employee is Hispanic	-0.001	-0.007**
	(0.001)	*(0.002)*
Employee is Asian	-0.001	-0.007**
	(0.001)	*(0.002)*
Current manager is Black	0.006‡	0.001
	(0.003)	*(0.003)*
Current manager is Hispanic	0.001	-0.003
	(0.002)	*(0.003)*
Current manager is Asian	-0.002	-0.007‡
	(0.004)	*(0.004)*
White employee * White manager		-0.007**
		(0.002)
Black employee * Black manager		0.005
		(0.005)
Hispanic employee * Hispanic manager		0.001
		(0.004)
Asian employee * Asian manager		0.005
		(0.006)
Observations	>50,000	>50,000
R-squared	0.03	0.03

Note: Sample is all employees who quit during the sample period. Estimates are from linear probability models with store fixed effects plus the additional control variables noted in Table 3. Robust standard errors in parentheses, adjusted for clustering on manager spell. ‡ Significant at 10 percent; * significant at 5 percent; ** significant at 1 percent.

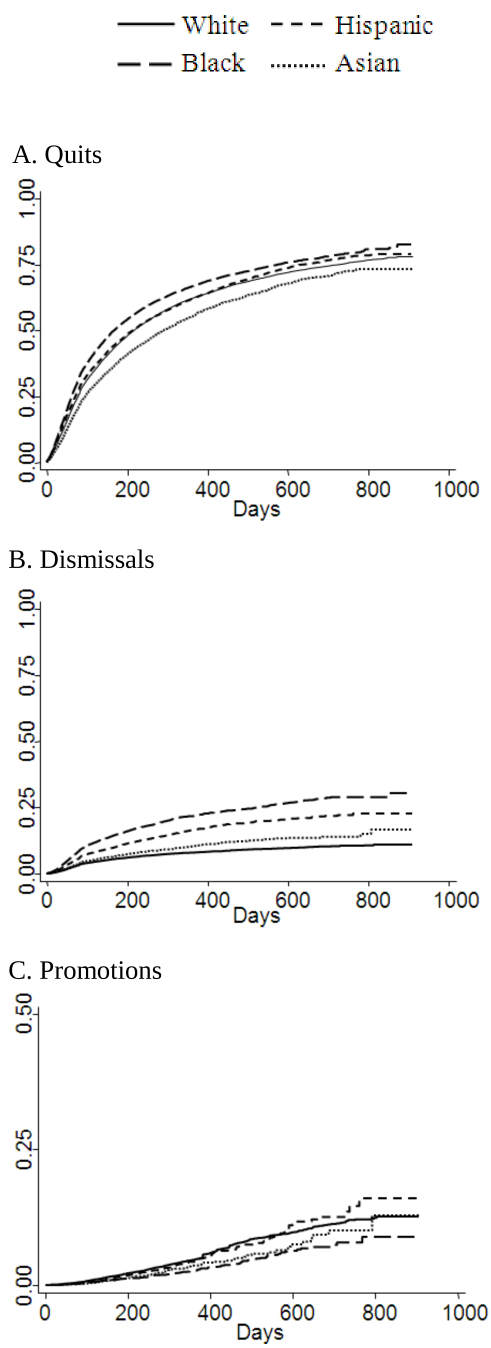

Figure 1
Kaplan-Meier Failure Estimates, by Employee Race