Learning to Discriminate on the Job

Alan Benson
University of Minnesota

Louis-Pierre Lepage
Stockholm and Queen's University

December 2021

Abstract

Using administrative records from a large national US retailer, we find managers learn to discriminate "on the job" as they gain experience hiring workers of different races. First, we find that positive and negative experiences with black hires seed the race of future hires, consistent with managers updating their beliefs about the productivity of these workers. Second, the degree of updating is larger for black than white workers. Third, because they decrease hiring, negatively-biased beliefs are slower to self-correct than positively-biased ones. Fourth, these dynamics, combined with the fact that black workers are already in a minority, yield particularly persistent declines in black hiring following managers' negative hiring experiences. Our results suggest that managers develop biased beliefs from endogenous learning about racial groups, which systematically disadvantages minority workers.

JEL Classifications: J71 (Discrimination); M50 (Personnel economics); D83 (Search, learning, information and knowledge, communication, belief, unawareness); J24 (Human capital, skills, occupational choice, labor productivity)

Keywords: Labor market discrimination, managers, employer learning, belief formation

We thank Martha Bailey, John Bound, Charlie Brown, David Card, Steve Lehrer, Matthew Lindquist, Conrad Miller, David Miller, José Montalbán Castilla, Rosina Rodriguez Olivera, Dan-Olof Rooth, Jeff Smith, and Aaron Sojourner for helpful discussions and feedback as well as seminar participants at CEA, EEA, SOLE, CIREQ, EALE, the University of Michigan, SOFI at Stockholm University, and the Online Seminar on the Economics of Discrimination and Disparities for their comments. A previous draft of this paper was circulated under the title "The Emergence of Hiring Discrimination."

Extensive and persistent racial disparities pervade the labor market. The full-time male black-white earnings ratio was 0.77 in 2010 with little progress since the 1990s, while the relative unemployment gap has remained approximately constant for decades (Lang and Lehmann, 2012). Substantial recent effort has sought to understand the root causes of discrimination and its impact on labor market disparities. In particular, recent studies have traced discrimination to biased beliefs, in the sense of incorrect group perceptions, among individual recruiters, managers, judges, physicians, and other influential decision-makers who can be enormously consequential in shaping economic outcomes (Reuben et al., 2014; Arnold et al., 2018; Bohren et al., 2019; Bordalo et al., 2019; Sarsons, 2019).

But where do individuals' biased beliefs come from? Most theories posit that discrimination is a product of market-level group differences, as in classical statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977), or a decision-maker's personal taste (Becker, 1957), leaving little room for incorrect beliefs. In contrast, recent work has considered biased beliefs arising from the stereotypical exaggeration of true group differences (Bordalo et al., 2016), implicit group associations (Bertrand et al., 2005; Rooth, 2010; Glover et al., 2017), or attribution bias (Sarsons, 2019). An alternative possibility is that biased beliefs emerge specifically from a decision-marker's immediate experience with their environment. This possibility would be most consistent with growing evidence across economics and finance that individual beliefs are largely governed by domain-specific personal experiences (Malmendier, 2021a,b). The possibility that managers update their beliefs based on their personal hiring experiences may help explain broader patterns of labor market discrimination and generate distinct policy implications (Lepage, 2021a). However, data limitations have stymied efforts to study the emergence of individuals' biased beliefs and discriminatory hiring in labor markets (Charles and Guryan, 2011; Guryan and Charles, 2013).[1] For instance, Census and audit study data typically lack information on individual hiring managers, inhibiting the ability to study the emergence and evolution of biased beliefs.

This paper seeks to overcome this challenge by using administrative data from the US operations of a large national retailer to examine how managers' past experience hiring workers of different races affects the race of their subsequent hires. The data include over

[1] A broader literature has documented heterogeneous effects by a manager's group membership, e.g. by race or ethnic background (e.g. Giuliano et al., 2009; Giuliano and Ransom, 2013; Åslund et al., 2014; Hjort, 2014; Glover et al., 2017; Berson et al., 2020) or gender (Cullen and Perez-Truglia, 2021; Ronchi and Smith, 2021)

1 million workers in permanent positions working under 27,000 store managers across 4,900 stores between 2009 and 2016. The data are particularly well suited to study the evolution of manager-level hiring discrimination: hiring is highly decentralized and at the discretion of department managers, who are free to draw upon their past experience to make hiring decisions. Because department managers hire for departments nested within stores, we are able to isolate manager effects from the effects of the job or store location. The data also afford relatively high power to study the evolution of hiring across a large set of managers; about half a percent of the stock of the US labor force was hired by the firm in this period. Workers in the retail-trade sector constitute about 10% of the US labor force and share similar barriers to economic mobility as other working class occupations (Bureau of Labor Statistics, 2021).

We begin by documenting substantial variation in the race of workers hired by managers that is not explained by location or jobs. Our main analysis concerns whether manager-level variation in hiring itself reflects dynamic belief formation "on the job." To do so, we consider a theoretical framework adapted from Lepage (2021a). Managers are initially uncertain about the performance distribution among applicants of different racial groups (e.g. white and black), particularly minority groups, about whom managers have less initial information and therefore weaker priors.[2] They update their beliefs as they hire minority workers, with positive experiences leading managers to increase their hiring of the group on expected productivity grounds. In contrast, negative experiences reduce future hiring, presenting fewer cases for managers to update their beliefs. Therefore, the model predicts that negatively-biased beliefs self-correct more slowly than positively-biased ones. The larger magnitude of belief updating and rarity of subsequent hiring following negative experiences with minority hires yield persistently lower hiring rates for these workers.

To operationalize positive and negative experiences, we use realized tenure among new hires of different races for permanent positions. Turnover at this firm (and in retail generally) is very high, as are the costs of recruiting, training, and ramping up new workers. Among workers for whom objective performance measures are available, turnover is also strongly associated with poorer performance and lower rates of promotion. Our main analyses

[2]The idea that employers have noisier information about minority workers is consistent with previous work in statistical discrimination (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009), but we explicitly consider the dynamic implications of noisier information for subsequent hiring and belief updating.

treat positive experiences constitute new hires who last at least 12 months, and negative experiences constitute new hires who are fired or quit within 3 months. Our main results follow.

First, positive past experiences with black workers increase a manager's propensity to hire from that group, whereas negative experiences decrease their propensity. We argue that, if beliefs were static, hiring for different groups would vary across managers, but not depend on whether previous individual hiring experiences were positive or negative. Instead, this result is consistent with the proposition that managers update their beliefs while on the job, and not with the proposition that discriminatory beliefs and behaviors are entirely fixed by the time they become managers.

Second, the effects of experience on future hiring are particularly pronounced for black workers. Given that the distribution of positive and negative experiences is similar across black and white workers at the firm, one interpretation is that managers have relatively weak priors about the performance distribution of black workers or that black workers are a minority, and as a result, similar information yields greater changes in beliefs. Behavioral biases, particularly stereotyping, may also amplify updating if managers generalize their experiences with minority workers (Allport et al., 1954).

Third, hiring is asymmetric across positive and negative experiences. Early negative experiences with black workers lead to persistent declines in black hiring, while the effects of early positive experiences are comparatively short-lived. Intuitively, negative experiences reduce future hiring, which slows the rate at which managers correct their negatively-biased beliefs. In contrast, positive experiences yield more future hiring and a larger sample of hires from which to correct their positively biased beliefs. As a result, managers on average systematically underestimate the performance of black workers at the firm.

Fourth, negative experiences with black workers have more persistent impacts than for white workers, following from greater updating (result 2) and naturally lower rates of minority hiring. Combined with the relative persistence of negatively-biased beliefs (result 3), this result implies that the initial "seeding" of hiring experiences with minorities can have substantial, persistent effects on hiring discrimination, namely by leading managers to underestimate the performance of black workers on average.

We investigate several potential alternative mechanisms through which managers' hiring histories could affect their subsequent hiring. One broad class of alternative explanations

attributes the race of future hires to the hiring history of the business unit rather than the manager. For instance, employee discrimination, referrals, selection on unobservables, and productive complementarities could also explain why the race of a department's past hires predicts future hires. We find that time and manager placebos as well as team composition variables fail to predict the race of future hires, whereas including terms for hiring experience of the "true" manager adds substantially to our ability to predict future hiring. Our placebo analyses therefore require alternative explanations to be manager-specific, such as stationary taste-based discrimination by the manager. However, stationarity is not consistent with our results, namely regarding how hiring evolves within an individual manager's tenure. We caution that our tests do not rule out that these other factors may be at play. Rather, our results are consistent with the proposition that the extent of discriminatory hiring depends on the individual manager, and that hiring within managers evolves in a way consistent with belief updating on the job. In other words, accounting for on-the-job hiring experience of individual managers adds to our ability to predict hiring discrimination.

Our results relate to the increasing interest in discrimination research that examines discriminatory beliefs and behaviors of individual decision-makers, which has largely set aside the question of where such beliefs originate and how they evolve. We contribute to a growing body of work on managers, particularly in decentralized organizations, having discretion in hiring which leaves room for individual biases (Hoffman et al., 2018; Berson et al., 2020; Bergman et al., 2020). Our results also relate to the literature specifically on belief formation, particularly evidence on "experience effects" (Malmendier and Nagel, 2011; Malmendier, 2021a,b).[3] These studies have found that decisions are largely governed by subjective beliefs formed over personal experiences, and not by group-level averages or expert assessments. Our results suggest discriminatory behaviors are also largely governed by experience effects that operate similarly in labor markets (Leung, 2018; Lepage, 2021b).

The form of discrimination we document is grounded in a lack of group-level information or learning, which differs from traditional theories of taste-based and statistical discrimination in how discrimination arises, evolves, and can be mitigated (Bohren et al., 2021; Hull, 2021; Lepage, 2021a). Corroborating such effects provides reasons for both

[3]Similarly, our findings suggest that when an employer hires a worker and observes their productivity, they not only update about the individual worker as in the classical employer learning literature (Farber and Gibbons, 1996; Altonji and Pierret, 2001; Lange, 2007; Arcidiacono et al., 2010; Kahn and Lange, 2014), but also their group.

pessimism and optimism for combating discrimination. On the one hand, minorities are inherently disadvantaged because negatively-biased beliefs about them are larger and more persistent, even without invoking behavioral biases, biased priors, productivity differentials, or fundamental prejudice. On the other, an important driver of hiring discrimination appears to be mistaken manager beliefs, suggesting room for new organizational practices like information-based policies and contact-based interventions (Paluck et al., 2019; Bergman et al., 2020).

The rest of the paper is organized as follows. In Section 1, we present our theoretical framework. In Section 2, we describe our data and institutional setting as well as characterize black hiring across managers at the firm. Section 3 presents our empirical approach, while Section 4 presents our main results on how previous experiences of managers with groups generate hiring discrimination. Section 5 presents additional results regarding employer learning from experience, while Section 6 investigates alternative interpretations of our main results. Lastly, Section 7 concludes.

1 Theoretical framework

We present a theoretical framework in which managers iteratively update their beliefs about the productivity of worker groups based on the observed productivity of their hires, and hire based on their (endogenous) beliefs. Our primary proposition is that manager beliefs about a worker group's productivity become biased, because they depend on observed hires which themselves depend on beliefs. Managers "learn" from a biased sample of worker productivity, and these biased beliefs are seeded by previous experiences with workers from the group.[4]

The model is adapted from Lepage (2021a), focusing on shorter-run individual hiring decisions of managers, from a finite exogenous set of applicants, and at an exogenous wage set by the firm. Our exposition considers managers who condition beliefs on race, but in principle, the model could be applied to any group characteristic, such as education or personality. Moreover, although managers in the model update their beliefs about the general

[4]This type of discrimination differs fundamentally from taste-based discrimination. While both generate outcome differentials that are not grounded in true group differentials, biased beliefs about productivity are caused by a lack or information or learning, rather than a fundamental prejudice. Taste-based discrimination is typically not expected to respond systematically to changing beliefs about productivity, leading to different dynamics and policy implications.

productivity of black workers, we emphasize that managers may only be updating their beliefs about the performance of hires given their location or job, and that they may value other aspects like a worker's likelihood of turnover. Although it's unclear whether managers generalize their experiences to the full universe of people of a given group, our analysis suggests that managers perceive their prior experiences as informative at least within the same parent company.

As is standard in statistical discrimination models, managers hire workers based on their expected productivity and, in the absence of perfect information on individuals, use group membership as a potentially relevant indicator of individual productivity. Unlike standard models of statistical discrimination, we do not assume that employers have complete information on group productivity or that their beliefs are confirmed in equilibrium, for example through endogenous worker responses (Coate and Loury, 1993). Rather, managers are initially uncertain about the productivity of groups, particularly minority groups about whose productivity they have noisier initial information (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009). Then, experiences with these workers play a disproportionate role in shaping belief updating and subsequent hiring.

Specifically, a manager is tasked with hiring the most productive worker from a candidate pool of two groups denoted by W and B. Our framework focuses on managers updating their beliefs about the minority group B, which we posit is the key driver of hiring discrimination. Managers take vacancies and entry wages for a position as given. The expected productivity of worker i from group B, x_{iB}, depends on a noisy signal of individual productivity s_{iB} observed prior to hiring and group membership. The individual signal is composed of the worker's productivity and an unbiased noise component: $s_{iB} = x_{iB} + \varepsilon_{iB}$ with $\varepsilon_{iB} \sim N(0, \sigma^2_{\varepsilon B})$. For example, it could include information from a resume, pre-employment test, or interview.

Worker productivity is normally distributed with mean μ_B and variance σ^2_B, such that $x_{iB} \sim N(\mu_B, \sigma^2_B)$. The productivity distribution of group B is static, but beliefs about it evolve heterogeneously over time. For expositional simplicity, we assume that managers know the variance in group productivity σ^2_B and the noisiness of individual signals $\sigma^2_{\varepsilon B}$.[5] As

[5]Employers learning about productivity variance or individual signal precision with experience affects the relative weight attached to individual signals versus group membership, but not the substantive implication that positive (negative) experiences increase (decrease) group B's expected productivity, and therefore the probability of hiring from the group in the future.

such, we focus on posterior beliefs about the mean productivity of group B, $\hat{\mu}_B$, and define $S_{Bmt} = \{x_{iBn} : i \text{ from } B \text{ is hired by } m \text{ at time } n\}_{n=1}^{t}$ as the information set about workers from group B available to manager m at time t.

The expected productivity of worker i from group B for manager m at time t is

$$P_{iBmt} = E[x_{iB}|s_{iB}, E[\hat{\mu}_B|S_{Bm,t-1}]] = \gamma_{Bmt} s_{iB} + (1 - \gamma_{Bmt}) E[\hat{\mu}_B|S_{Bm,t-1}]$$

where $\gamma_{Bmt} = \frac{\sigma_B^2 + Var[\hat{\mu}_B|S_{Bm,t-1}]}{\sigma_B^2 + Var[\hat{\mu}_B|S_{Bm,t-1}] + \sigma_{\varepsilon B}^2}$.[6] The manager hires the worker with the highest expected productivity out of the set of applicants A_t, with fraction F_{Bt} from group B.[7] That is, worker i from group B is hired at time t if $P_{iBmt} > P_{i'g'mt}$ for all $i' \in A_t$ from group g', and for $g' \in \{W, B\}$. When manager m begins hiring, they have a prior belief distribution about group B's mean productivity $\hat{\mu}_B|S_{Bm0}$. If they don't hire from group B at time t, they don't update their beliefs and $\hat{\mu}_B|S_{Bm,t-1} \overset{d}{=} \hat{\mu}_B|S_{Bmt}$. If they hire from group B at time t, managers observe their hire's productivity x_{iB}, which is not randomly drawn because managers select workers based on expected productivity.

Accordingly, managers first form an expectation about x_{iB} given that worker i has the highest expected productivity out of all applicants

$$E[x_{iB}|P_{iBmt} > P_{i'g'mt} \ \forall i' \in A_t, g' \in \{W, B\}].$$

Second, managers update their beliefs from $\hat{\mu}_B|S_{Bm,t-1}$ to $\hat{\mu}_B|S_{Bmt}$. Posterior variance monotonically decreases with additional hires, while the direction of updating about the mean depends on the discrepancy between the hire's expected and observed productivity

$$E[x_{iB}|P_{iBmt} > P_{i'g'mt} \ \forall i' \in A_t, g' \in \{W, B\}] - x_{iB}. \tag{1}$$

If realized productivity is above (below) expectation, denoted as a positive (negative) hiring experience, managers update their beliefs based on equation (1) with upwards updating when the difference is negative and downwards updating otherwise. A positive

[6] Employers know σ_B^2 for a given mean, but uncertainty about the mean introduces additional variance in expected productivity $Var[\hat{\mu}_B|S_{Bm,t-1}]$.

[7] As modeled in Lepage (2021a), managers should also value learning about group B to improve subsequent hiring decisions. We abstract from this feature since it does not change the direction of belief updating and that of its impact on hiring and learning across groups.

(negative) experience with group B increases (decreases) $E[\hat{\mu}_B|S_{Bmt}]$ and $P_{iBm,t+1}$ relative to group W, increasing (decreasing) the probability that a group B worker is hired in period $t+1$.

The rate of hiring also drives the speed of learning. Contrasting positive and negative hiring experiences, a positive experience, through increasing the probability of subsequent hiring from the group, also increases the probability of observing signals about the group's productivity at time $t+1$ and beyond. Depending on the manager's prior, these additional signals may partially undo the impact of the positive experience through reversion to the mean and in any case lead to more accurate beliefs. In contrast, a negative experience, through decreasing the probability of subsequent hiring from the group, also decreases the probability of observing signals about the group's productivity at time $t+1$ and beyond. Regardless of prior beliefs, the impact may be more persistent, because it decreases learning.

In addition, since workers from group B are baseline more infrequently hired, manager experiences with them may also have more persistent impacts since learning about their productivity is slower and it is likely easier (harder) for managers to avoid hiring (seek out) groups who constitute a minority of applicants following negative (positive) experiences.

Disproportionate updating, combined with more persistent impacts of negative experiences on subsequent hiring, predict that hiring experiences systematically decrease relative group B hiring across managers. To summarize, our framework yields four key predictions regarding managers' beliefs updating

1. Managers update their beliefs of group B productivity from their hiring experiences.

2. Managers disproportionately update their beliefs of group B, about whom they have weaker priors.

3. Managers' negatively-biased beliefs about group B productivity correct more slowly than positive ones.

4. Managers' negatively-biased beliefs about group B productivity correct more slowly than those about group A productivity.

These predictions don't rely on prior bias or prejudice, and they are robust to considering deviations from Bayesian updating that do not impact the direction of belief updating. They

also do not rely on whether worker groups have true productivity differences: whatever the true productivity of minority groups is, employers will underestimate it on average. They differ from predictions of other, typically static, types of bias established in the literature such as inattention discrimination, implicit biases, attribution bias, negative stereotypes, and self-fulfilling prophecies, and are robust to considering them. In fact, these other potential sources of bias are largely complementary since they would generally predict lower group B hiring by managers, slowing down learning and amplifying biased beliefs that arise from managers' experience in our framework.

Since we do not observe beliefs directly, we implement our tests under the condition that belief updating is reflected through hiring decisions. This yields the following four testable predictions regarding how individual managers hire black and white workers at the firm

1. Within managers, positive/negative experiences with black workers positively/negatively affect the share of future hires who are black.

2. Within managers, positive/negative experiences with black workers disproportionately affect the share of future hires who are black, compared to the manager's positive/negative experiences with white workers.

3. Within managers, negative experiences with black workers have a more persistent impact on subsequent hiring than positive experiences.

4. Within managers, negative experiences with black workers lead to more persistent declines on subsequent hiring than negative experiences with white workers.

We observe performance measures, but not their discrepancy with a manager's expectation, so we use different measures relative to other workers at the firm, which should inform expectations.

2 Setting

Our data consist of monthly longitudinal administrative records on workers and managers from the US operations of a large national retailer between February 2009 and October

2016. For each worker and manager, we observe tenure, demographics, job, department, and location. We also observe employment termination including dismissals, quits, and layoffs along with stated reasons for dismissals and quits. Each store is led by one store manager and a set of department managers who hire for their respective department (on average 5-6 per store), allowing us to study hiring decisions of each department manager over time. We restrict our sample to new hires into permanent non-managerial retail positions, as these are presumably the most consequential for the manager and positions for which tenure can be used as a measure of the worker's performance (Autor and Scarborough (2008)). Excluding transfers or returning workers allows us to concentrate on hires that were chosen by the manager specifically in the given hiring period and are therefore likely more salient.

We focus on white and black workers because these are the two largest racial categories in our data, which makes it most feasible to estimate managers' evolving hiring behavior.[8] Summary statistics on workers and managers are presented in Panel A of Table 1. In particular, black and white workers account for nearly 80% of hires, slightly more than half of workers are female, and managers on average hire 5-6 workers per year.

To study how managers' previous experiences hiring workers influence their subsequent hiring decisions, we consider both a time panel which captures previous experiences in months leading up to a manager's current hiring decision as well as a manager-level panel in which one observation corresponds to a month in which a manager hires at least one black or white worker, which we refer to as a hiring event. Our baseline analysis restricts our sample to new managers who began hiring at the firm for the first time during our sample period, although we also present analyses including managers who hired workers prior to entering their current position. On average, managers hire workers approximately every two and a half months. We observe 60,096 hiring events (46% of all manager-months) with an average of 2.3 workers per hiring event (0.75 black, 1.55 white). One motivation for organizing the data by hiring event is to focus on belief updating from managers' own hiring experiences, which is likely concentrated in periods managers actually hire workers and arises irregularly across time.[9]

[8]Hispanics are treated as a separate category in the data and corresponding analyses are presented in Appendix D. Evidence on differentials between Hispanic and white workers is more mitigated and harder to interpret since the firm does not distinguish between race and ethnicity.

[9]In these specifications, we control for the number of hires in a given hiring event as well as the time between hiring events to account for the possibility that negative and positive experiences affect the time in between hiring events. Ultimately, results shown below indicate that managers' updating behavior from

Staffing levels for permanent positions in each store are determined by forecasts made by the firm's headquarters. When a manager wishes to fill a vacancy, the manager would typically begin by requesting a shortlist of candidates from the location's designated HR representative. The manager can then interview candidates and make offers. Workers in most entry-level jobs (e.g. cashiers, sales associates, and material handlers) are provided one week of formal online skills training and a week of job shadowing before moving to regular status. Most entry-level positions are filled from "evergreen requisitions," meaning candidates can apply at any time and may be called for an interview as needed. Positions may also be filled by department managers who conduct informal or spot interviews with candidates prior to submitting a formal application, and then notify their HR representative of their interview performance. Anecdotal evidence from store managers we interviewed indicates that it would be rare for a manager to have prior familiarity with an applicant, though there is no formal process for tracking referrals and we do not observe these instances directly.

Applicants take an online screening test that classifies them into three recommendation tiers. Department managers observe screening test results, but are free to deviate from the algorithmic recommendation when selecting applicants to interview. Department managers are trained in behavioral interview techniques. For instance, a customer service applicant may be asked how they might respond to a hypothetical scenario, like dealing with a difficult customer, or to discuss an instance they confronted a problem and were proud of the solution they offered (Dipboye et al., 2004).

Turnover at the firm is high, in line with the retail sector more generally which has 50% greater turnover than the US average (JOLTS, 2017). High turnover provides valuable variation in hiring of workers within managers even within a limited time horizon, allowing us to better study adjustments that managers make with hiring experience. Survey evidence indicates that the average cost of hiring and training a replacement retail worker is $3,328 or around 10 weeks of worker salary (Boushey and Glynn, 2012). Accordingly, retaining workers who have been hired for permanent positions is one of the most important aspects of department managers' jobs.

Figure 1 shows that nearly 90% of workers no longer work in the position for which they were hired at a given store after one year. Most turnover reflects dismissals and quits,

their hiring experiences is qualitatively similar whether considering a time or a manager-event panel.

especially in the first three months of employment, with around 68% versus 52% after one year. Other sources of turnover for a department manager are transfers across departments and stores, layoffs, and promotions/disability/death, which account respectively for 18%, 25%, and 5% of turnover at the 12 month mark. After the first year of employment, the turnover rate falls substantially and remains fairly stable below 2% per month.

There is large variation in black hiring across managers at the firm. Figure 2 shows the share of black and white workers hired over our sample by each manager. The mean share of black workers hired is 20%, the median is 8%, and more than a quarter of managers hire no black workers. The mean share of white workers hired is 56% and the median is 59%. In Appendix B, we present evidence of substantial cross-manager variation in the share of black hires, which persists after controlling for the manager race as well as store, department, and job effects. A substantial share of residual variation in black hiring is explained by manager fixed effects, implying that individual manager effects play a substantial role in determining the race of hires.

3 Belief updating from experience

Our empirical analysis examines whether manager-level variation in hiring can be explained by managers' updating from their own prior experiences.

Our main specification investigates how experiences with black and white workers in previous hiring events affect the share of black hiring in the current hiring event. We estimate the following model

$$FB_{emlt} = \beta_1 E\bar{X}P_{B,e-1} + \beta_2 E\bar{X}P_{W,e-1} + X_{emlt}\zeta + \theta_t + \lambda_l + \gamma_m + \varepsilon_{gemlt} \tag{2}$$

where the dependent variable is the fraction of black workers hired in hiring event e by manager m in location l at time t. The primary coefficient of interest is β_1, capturing the impact of more positive or negative average hiring experiences with black workers up to hiring event $e-1$. Similarly, $E\bar{X}P_{W,e-1}$ is a measure of more positive or negative average hiring experiences with white workers up to hiring event $e-1$. X_{emlt} includes the fraction of full-time workers, fraction female, average age, total number of hires, number of previous hiring events, time since last hiring event, yearly state unemployment, and yearly state college

attainment. θ_t, λ_l, and γ_m represent month and year, store, and manager fixed effects. Store fixed effects account for differences between applicant pools, local markets, and store-level characteristics faced by the manager, among other factors. Manager fixed effects account for time-invariant manager differences that may affect their willingness or ability to hire applicants of different races. Standard errors are clustered at the manager level.

This specification can be used to test our theoretical framework's key predictions by examining coefficients on $E\bar{X}P_{B,e-1}$ and $E\bar{X}P_{W,e-1}$. Intuitively, near-zero estimates suggest that heterogeneity in the race of manager's hires can be explained by factors relating to the hiring context and (potentially unobserved time-invariant) manager characteristics. In contrast, if past experience predicts future hiring net of these other factors, we would interpret that as evidence that manager group perceptions are not fully formed or immutable by the time they begin hiring. Rather, positive and negative experiences with a group would appear to affect manager's beliefs and therefore their propensity to hire from the group.

3.1 Experience measures

Because our analysis considers hiring decisions as a function of idiosyncratically positive or negative previous experiences, we must distinguish hiring events as either positive or negative versus a manager's expectations for both worker races. We use two main approaches to classify experiences as positive or negative.

First, organizing the data as a time panel, we compare tenure achieved by a manager's hires to expected tenure at the firm. Specifically, we compute deviations in turnover rates by race and tenure at the level of the manager's subordinates from average turnover rates at the firm, cumulatively for every month leading to a given hiring event.[10] The cumulative average of these deviations indicates how a manager's previous hires from each racial group were more or less likely to achieve a given level of tenure than expected.

Second, organizing the data by hiring event, we identify particularly negative and positive experiences of a manager, focusing on specific experiences likely to be most salient. For negative experiences, for each hiring event, we calculate the share of each race that was fired or quit in the first 3 months of employment. As shown in Figure 1, the first 3 months

[10]Results are qualitatively and quantitatively similar comparing to race-specific turnover rates or average turnover rates across racial groups.

represent a key period after which the rate of dismissal and quit decreases substantively.[11] Workers hired into permanent positions who leave or are terminated within the first 3 months account for around a quarter of hires. They are also very costly: these workers must be hired, paid during a training period, provided time to develop tacit skills and a familiarity with the store's protocols, procedures, and products. Workers who depart after short tenures are also costly in terms of their opportunity cost: they filled a spot that could have otherwise been filled by a successful hire.

For positive experiences, for each hiring event, we calculate the share of each race that achieves at least one year of tenure in the position for which they were hired. As shown in Figure 1, after 12 months, the likelihood of a position separation in any given month is fairly low and stable. Long tenure suggests a successful hire and sufficiently good match between the worker and the position. It also reflects a stronger measure of worker performance than using the share of workers that has not quit or been fired after 12 months, since we may be concerned that poorer hires could also be transferred or laid off. Approximately 10% of new hires achieve tenure of at least one year in their position. Alternatively, to keep experience measures symmetric across positive and negative experiences, we can define positive (negative) experiences as the previous hires of a manager being in the top (bottom) quartile of deviations from expected tenure at the firm, which as shown in Table C yields similar results. Given the potentially forward-looking nature of our measures, we exclude workers hired in the last 3 months (1 year) of our sample for negative (positive) experiences.

In principle, even though managers likely directly value tenure in a position, we would ideally also have a direct measure of productivity that the manager observes. In that case, we would test whether future hiring depends on productivity rather than tenure signals for a race. Still, standard theories feature a positive relationship between productivity and tenure (such as job search models), and empirical work has established such a relationship with regularity.[12] Appendix A presents additional evidence of this relationship at the specific firm we study.

By using tenure to classify positive and negative hiring experiences, we are not

[11] Using measures of the performance of black workers relative to white workers or relative to workers in the store's CBSA has limited impact on the results, as shown in Columns 3-4 and 11-12 of Appendix Table C.

[12] See for example Bycio et al. (1990), Williams and Livingstone (1994), Huang et al. (2006), and Zimmerman and Darnold (2009).

asserting that objective worker performance is the only factor affecting hiring and retention. Differences in average tenure across race, although small as shown below, could be explained by differences in performance, but also differences in average discrimination. Such discrimination could take several possible forms: managers may require higher performance for minorities as a condition for continued employment, managers may put greater effort in training non-minorities, or minorities may shirk under biased managers Berson et al. (2020); Glover et al. (2017). However, our goal is not to assess differences in productivity or hiring in the absence of bias, but rather to examine whether bias evolves on the job based on an employer's potentially subjective experiences. Further, we show that static or pre-existing bias, without involving updating by managers, provides a poor alternative explanation for the patterns of hiring discrimination we document.

Panel B of Table 1 shows summary statistics for our performance measures. Compared to white hires, black hires have a slightly higher 3-month quit or dismissal rate and slightly lower 12-month retention rate. Importantly, most variation in tenure is not explained by race. For instance, both black and white workers have the same median tenure of 4 months. Because managers only hire an average of 5-6 workers per year, and fewer of each race, they are left with relatively few personal observations from which to update beliefs. We argue that this raises the possibility that biased beliefs formed by unlucky initial draws, particularly with black workers, could take a long time to self-correct, especially if they endogenously reduce hiring of these workers.

4 Main results

4.1 The effect of prior experiences on future hiring

Our first proposition is that a manager's previous experiences hiring black workers will influence their subsequent hiring of these workers.

Table 2 and Table 3 estimate the effect of positive and negative previous hiring experiences on the race of subsequent hires. The outcome variable corresponds to the share of hires that are black, but since the sample is restricted to black and white workers, estimates for the fraction of white hires are the same magnitude but opposite sign. The independent variables capture the cumulative impact of previous experiences with each race. Estimated impacts in

percentages are proportionally larger for black than white hiring given that they constitute a minority of workers, approximately 50% larger, indicating that hiring experiences play a disproportionately large role in black hiring.

Table 2 presents results for the measure of expected tenure by race, with more negative (positive) experiences resulting in a more negative (positive) measure of deviation from expected tenure. Column 1 restricts the sample to managers who started hiring at the firm during our sample period, while Column 2 includes all managers with a hiring spell that began during our sample period. The results show that a higher expected tenure for black workers based on previous experiences leads to statistically significantly more hiring of these workers, with a one standard deviation increase in expected months of tenure corresponding to an increase of approximately 5% in relative hiring of black workers. The estimates for expected tenure with white workers are negative as expected, but they are substantively smaller and statistically non-significant.

The first two columns of Table 3 present the results of negative experiences specifically. They indicate that managers statistically significantly decrease their hiring of black workers by an estimated 6% for a one standard deviation increase in the fraction of previous black hires that were dismissed or quit within 3 months. Estimates across columns suggest that these impacts do not only affect new managers. Estimated impacts for experiences with white workers indicate a substantially smaller but still statistically significant increase of approximately 3% in black hiring, when accounting for the higher standard deviation of experience measures with black workers. Columns 1 and 2 of Appendix Table C show that both quits and fires have a similar negative impact on subsequent hiring, consistent with managers wanting to avoid both and inconsistent with some of our results being driven by a differential share of fires versus quits across racial groups (Autor and Scarborough, 2008).[13]

Columns 3 and 4 of Table 3 present estimates of the impact of positive previous experiences. Managers statistically significantly increase their hiring of black workers by 4% for a one standard-deviation increase in the fraction of previous black hires who reached at least one year of tenure in their position. Estimated impacts for white workers are smaller and statistically non-significant.

[13]Updating from negative experiences across groups is similar in periods of high versus low labor market tightness, which may influence the relative quality of outside options across race, also suggesting that our results are not driven by differential turnover reasons across race.

Overall, the key takeaway is that across specifications, cumulative previous experiences with black workers have a substantial impact on subsequent hiring decisions that is not restricted to new managers. Appendix Appendix D also shows qualitatively similar results when restricting to female workers or black managers, suggesting that previous experiences with both black male and female workers affect the race of future hires and that both black and white managers respond to their previous experiences with black workers.

4.2 Greater updating for black workers

Our second proposition is that experiences of managers with black workers disproportionately affect their future hiring. Across Table 2 and Table 3, estimated coefficients on previous experiences with black hires are larger than those with white hires, for both positive and negative experiences, generally by over 30%, and are all statistically significant. Further, statistical tests reject the null hypothesis that impacts of experiences with black and white workers are equal but of opposite sign in most specifications, as shown by the p-values presented in the tables. Overall, evidence of updating is weaker and somewhat inconsistent regarding previous experiences with white workers, although coefficients are of the hypothesized sign. This suggests a substantial updating gap across races, which is key for generating discrimination.

We interpret this result in light of our theoretical exposition, which adopts a simple Bayesian approach: managers update their beliefs more for black workers because they have weaker priors. Weaker priors about black workers on average can arise naturally from the group being a minority and therefore being less information about them in the labor market. However, greater updating for black workers is also consistent with a set of behavioral biases. Psychologists assert that whiteness is largely invisible; Blacks are judged together while whites are assessed as individuals or along nonracial categories (Sue, 2006). A literature on stereotype formation suggests that managers may form associations about minority groups from relatively little information, yielding similar updating disparity (Allport et al., 1954). Since managers update more following both positive and negative experiences with black workers, our results do not appear driven by attribution bias (Sarsons, 2019), although we cannot rule out that similar biases also affect updating.[14]

[14]Similarly, if managers, potentially subjectively, expect that negative (positive) hiring experiences are

4.3 Persistence of early negative experiences

Our third proposition is that a manager's previous experiences with black workers may indirectly affect their subsequent belief updating by affecting hiring decisions. Accordingly, negative experiences, especially early on in a manager's career, may have persistent impacts on subsequent hiring by generating negatively-biased beliefs about the performance of black workers which decreases subsequent hiring and therefore learning. In contrast, the impact of early positive experiences may be more transitory if they generate positively-biased beliefs which correct themselves by leading to more hiring and learning.

We estimate a similar regression model to equation (3), replacing the measure of cumulative average experience with measures of experience with black and white workers hired in the manager's first hiring event with these workers and the average of their first three hiring events.

Figure 3 traces the persistence of managers' first hiring experience on their subsequent hiring decisions. Managers whose first experience with black workers was negative, such that their hire(s) were fired or quit within 3 months, reduced their subsequent share of black hires over their next 7-9 hiring events, corresponding on average to a 1.5-2 year period. The decrease is expectedly strongest in earlier hiring events and corresponds to a reduction of around 14% for the first 3 events and 5% for events 4-6. In sharp contrast, there is little evidence of persistent impacts on hiring following from a first positive experience with black workers or either a first positive or first negative experience with white workers.

Similar results are shown in Table 4. Estimates from the first column indicate a statistically significant 2% decrease in black hiring in the current hiring event for a one standard deviation increase in the fraction of the first black hire(s) that were fired or quit within 3 months. Estimates from the second column indicate a statistically significant decrease of 4% in subsequent black hiring for a one standard deviation increase in our measure of negative experience with black hires in the manager's first three hiring events.[15] These results thus suggest that early negative experiences with black workers impact subsequent hiring over our entire sample period. Columns 4-5 include hiring spells from more experienced

more (less) likely to arise with black workers, then they should update less (more) following such an experience, also in contradiction to our findings.

[15] Employers who hired fewer than three groups of black or white workers at the time of hiring event e are excluded.

managers which hired workers before the start of our sample, with smaller coefficients suggesting that negative experiences at the beginning of a manager's career, rather than at the beginning of their current hiring spell, are particularly impactful.

Columns 1-2 and 4-5 of Table 5 also show smaller and generally statistically non-significant impacts for early positive experiences, highlighting a different persistence pattern from early negative experiences with black workers in particular.[16]

4.4 Systematic declines in black hiring

Our final proposition follows from results established in the previous subsections: (1) managers' hiring of black workers is disproportionately driven by their previous experiences with those workers, (2) early negative experiences with black workers in particular lead to persistent declines in hiring of the group. In contrast, impacts of early negative experiences with white worker(s) are smaller and statistically non-significant as shown in Table 4. We can also reject that they are of the same magnitude but opposite sign to those on early negative experiences with black workers. Along with the fact that the arrival rate of black workers may be relatively low even in the absence of biased beliefs due to their minority status, this implies that managers persistently reduce their relative hiring of black workers following negative experiences, consistent with them developing persistent negatively-biased beliefs. Taken together, our results imply that hiring responses of managers following their experiences with workers systematically decrease the relative hiring of black workers.

5 Additional evidence on learning

5.1 Updating over a manager's hiring history

First, we investigate how a manager's most recent experience affects their current hiring. Recent experiences may be salient to managers, for example if the hiring context changes over time or due to recency bias (Agarwal et al., 2008; Gallagher, 2014; Erev and Haruvy,

[16]These results suggest that the subset of managers who hire black workers had roughly unbiased priors about their performance. If they systematically underestimated their mean performance, then negative experiences may have had a more muted impact on subsequent hiring and positive experiences may have lead to persistent increases.

2016). Columns 3 and 6 of Table 4 indicate a statistically significant decrease of around 5% in current hiring of black workers for a one standard deviation increase in the fraction of the latest black hire(s) that quit or were fired within 3 months. Estimates for white workers are approximately 30% smaller, but also statistically significant. Columns 3 and 6 of Table 5 present results for positive experiences corresponding to a 2% increase in black hiring for a one standard deviation increase in our positive experience measure for the latest black hires, and a 1.5% decrease for white hires. Since the most recent experience with both races appears salient, including for more experienced managers, the persistence of the impact of previous experiences with black workers appears to play a key role in generating sustained disparate impacts.

Second, since both early and recent experiences seem to play a role in shaping hiring, we break down how the impact of negative experiences evolves throughout a manager's hiring career. In particular, if managers learn with experience and become more confident about their beliefs regarding the performance of black workers, then we should observe a decreasing impact of negative experiences over their hiring history. Table 6 presents estimates of the cumulative impact of previous negative experiences with black workers on current black hiring, separating each manager's hiring events over our sample period into three chronological terciles. We compute the same measure of cumulative previous experiences as in Table 3, but separately within each tercile to see how much experiences in each tercile affect black hiring. We define terciles rather than pooling specific ranges of hiring events given substantial heterogeneity across managers in the number and timing of hiring events. Still, the results are qualitatively similar if we separate hiring events by whether a manager is in their first, second, or third and above years of hiring at the firm.

The results highlight that the impact of negative experiences that happened in the early, middle, and late segments of a manager's hiring history all affect black hiring. They also highlight that the impact of negative experiences decreases with a manager's hiring experience: the relationship between experiences and hiring is strongest early in a manager's hiring career and weakest in the last tercile of a manager's hiring history that we observe in our sample period. This pattern is consistent with learning decreasing variance in posterior beliefs and therefore adjustments in hiring behavior as employers gain experience.

Overall, results are consistent with managers putting substantial weight on their most recent experience, but within a broader learning process through which the weight they put

on their experiences decreases over their career.

5.2 Positive selection among workers who overcome bias

The theoretical framework proposes that managers condition their beliefs both on a worker's individual signal and their group. As a result, managers who draw low-performers from the pool of black applicants will be biased in their beliefs about that group, and set a higher threshold for a black worker's individual signal when making hiring decisions. Therefore, following a manager's negative hiring experiences with black workers, individuals of that group who are hired anyway should be less likely to yield a negative experience. Conversely, following a positive experience, the manager should lower the hiring bar for black workers, decreasing the probability of future positive experiences.

Consistent with this proposition, Table 7 shows that hiring experiences are negatively serially correlated. Estimates correspond to an approximate 30% decrease in the probability of a positive or negative experience for a one standard deviation increase in the cumulative fraction of previous positive or negative experiences. These results are inconsistent with reversion to the mean, since the quality of a current hire should be independent of the quality of previous hires in the absence of an additional mechanism operating at the manager or department level. They also imply that our measures of previous experiences are endogenous to an employer's beliefs in a specific way which may mitigate the persistence of negatively-biased beliefs: they decrease the likelihood of a subsequent hire, but increase its expected quality conditional on hiring.

5.3 Other person-specific sources of belief updating

Our main proposition is that managers' beliefs about worker groups are formed on the job from their hiring experiences, rather than being fully formed by the time they enter the job. In this section, we present additional evidence on how managers update their beliefs from different information sources, namely experiences of their colleagues, their own experiences at previous stores, their experiences with workers they "inherit" when joining a new department, and negative experiences unlikely to result from a bad match between the manager and the worker specifically.

Our first analysis examines the possibility that managers also update their beliefs from

the store's hires, and not just from their own hires within their department. In principle, the ability to communicate and learn from peers would accelerate the rate at which biased beliefs are corrected. However, even in this setting, where other same-store managers' experiences are fairly easy to observe, the cumulative average experiences of other managers at the same store have little impact on a manager's hiring after accounting for the manager's own experiences (see Columns 10 and 17 of Appendix Table C).

Our second analysis examines whether incoming managers appear to update their beliefs based on their hiring experiences in previous management roles at the firm. This analysis places considerable demands on the data, as we must restrict the analysis to 790 hiring events of new incoming managers for whom we also observe the outcomes of the manager's pre-move hires. It yields similar point estimates as our main analysis, though errors are outside conventional testing thresholds (see Column 9 of Appendix Table C).

Our third analysis examines whether incoming managers form beliefs based on the incumbent workers they "inherit," rather than just the ones they hire. To do so, we restrict the analysis to new incoming managers. We find that a more negative experience with black workers inherited by a new manager decreases subsequent hiring of the group by the manager, similarly to their first experience with their own hires, though this estimate is only statistically significant at the 10% level (see Column 8 of Appendix Table C).

Our fourth analysis examines whether negative experiences less likely to be endogenous to a manager's behavior also decrease hiring, looking at specific reasons for separation.[17] We find that these similarly decrease hiring, consistent with hiring decreases being driven by managers' updating beliefs about groups rather than reflecting their own behavior or their learning about their own managing ability.

Put together, our analyses suggest that a manager's belief formation is specific to their own context and somewhat portable across jobs, at least in a medium time horizon, though the time span of our data is limited (5 years). This sharp degree of specificity largely reflects the evidence on experience effects in other contexts (Malmendier, 2021a).

[17]These include dissatisfaction with pay, compensation, or benefits, which are not controlled by the department manager, as well as worker integrity and illegal or unethical behavior.

6 Alternative mechanisms

Our study's main claim is that managers develop hiring biases which are not static, but depend on their experience with worker groups. Thus far, we have found that hiring experiences with black workers do impact future hiring, and that early negative experiences in particular persistently reduce the rate at which managers hire black workers.

When evaluating threats to validity, one set of concerns regards factors correlated with the manager, particularly the manager's team. For instance, if workers refer prospective job candidates of their own race, then idiosyncratic differences in race within a team will be correlated with the race of subsequent hires. Similarly, if workers have a preference for same-race departmental colleagues ("employee discrimination"), then current racial composition will affect future racial composition. These processes, among others, would yield serial correlation in the race of a manager's hires that is not based on any factor specific to the manager or the manager updating their beliefs.

Our strategy for evaluating alternative mechanisms applies our prior findings, robustness checks, and falsification tests to establish that the effects we identify are specific to the manager and the timing of hires (and thereby the formation of beliefs) within their tenure.

First, our results are specifically attributable to individual manager effects and therefore inconsistent with any mechanism that operates through workers and customers like worker referrals, worker complementarities, and co-worker or customer discrimination. Table 6 highlights that the impact of negative experiences with black workers in a given department varies specifically with a manager's experience with worker groups. These results effectively act as a time placebo, showing that the largest impacts of negative experiences with black workers are at the beginning of a manager's career and decrease with hiring experience. Table C presents another time placebo investigating the impact of a more negative experience with black workers at a department in the month before a manager begins in their position, which has little impact on the manager's hiring, unlike their own first experience. Similarly, Appendix Table C highlights that there is but a negligible negative relationship (2% decrease for a one standard deviation increase) between the existing fraction of black workers in a department and the fraction of black workers hired in a given event. That is, a higher (lower) fraction of black workers, by itself, is not associated with a subsequent increase (decrease) in black hiring by a manager. Moreover, several alternative mechanisms predict

that positive and negative hiring experiences should be positively serially correlated, for example if high (low) productivity workers tend to refer other high (low) productivity workers from racially-homogeneous networks (Montgomery, 1991; Burks et al., 2015).[18] Yet, Table 7 shows that they are in fact negatively serially correlated, consistent with managers updating their beliefs about worker group productivity and adjusting their hiring thresholds accordingly. Regarding the specific hiring patterns we document, these alternative mechanisms at best provide little rationale for the relative persistence of early negative experiences with black workers in particular, rather than early positive experiences with black workers or early experiences with white workers, and at worse are inconsistent with it. Lastly, since departments are staffed with relatively few workers at any given time and draw from a fairly homogeneous pool of workers and customers, these alternative mechanisms likely play a larger role at the store level than at the department level.

A second set of concerns is that the effects we identify are specific to the hiring manager, but do not arise from them updating beliefs about the performance of worker groups from experience. Taste-based discrimination or other sources of bias could also create time persistence in hiring patterns. In the presence of prejudice, negative (positive) experiences could reflect a bad (good) working climate for minority workers which translates to less (more) hiring. Differently put, even if it is the manager's perception that is of interest for our purpose, rather than the worker's objective performance, the subjective assessment of a manager as to what constitutes a positive or negative experience may itself be biased and vary across groups. Pre-existing manager bias against a worker group could then affect both their expected tenure, for example by affecting how they are treated or evaluated by the manager, and the likelihood of the manager hiring from the group. Such bias could also generate what appears to be a persistent impact of negative experiences on subsequent hiring. Several of our results are inconsistent with this class of alternative explanations, particularly those based on static pre-existing biases or prejudice.

First, several of our specifications include manager fixed effects to account for time-invariant differences in group hiring across managers. Second, if the nature of hiring experiences is driven by underlying bias which correlates with subsequent hiring, then

[18] Previous work suggests that black workers are proportionally less likely to be hired through referral (Kirnan et al., 1989; Taber and Hendricks, 2003), inconsistent with the larger hiring responses that we document following experiences with black workers.

both negative and positive experiences should have a similar persistence, contrary to our results. Importantly, positive and negative experiences would also have to be positively serially correlated, in contrast to evidence presented above. Third, these alternatives provide little rationale for the decreasing impact of hiring experiences over a manager's career. Fourth, it seems unlikely that negative experiences with white workers reflect hostility against white workers in predominantly white departments with white managers. Yet, we find that experiences with white workers, at least in recent hiring events, do impact hiring. Similarly, coefficients from Table D suggest if anything *stronger* responses by black managers following experiences with black workers (an 11% decrease in black hiring for a one standard deviation increase in the fraction of previous black hires who quit or were fired within 3 months). Fifth, Table 1 shows that the rate of negative experiences with black and white workers at the firm is in fact fairly similar, indicating that much of the difference appears to lie in how managers respond to their experiences with these groups. Sixth, Appendix C shows that experiences with black workers who were fired or quit for reasons unlikely to be related to the manager's behavior also decrease black hiring, suggesting that evolving group perceptions play a role.[19]

To be clear, we do not interpret our results as indicating that managers are not biased against black workers in ways beyond that which arises through learning about their performance from experience. In fact, insofar as other sources of bias such as that documented in Glover et al. (2017) may arise from previous interactions with groups, they are largely complementary and consistent with our primary proposition that hiring experiences create group associations which lead to self-sustaining discriminatory behavior. Rather, we interpret our results as indicating that time-invariant, pre-determined biases provide a poor alternative explanation for the specific novel patterns of hiring we document, and that considering biased belief which evolve based on managers' experiences, in particular through learning, greatly helps rationalize the hiring discrimination we observe. This is especially true since we focus on dynamic biased beliefs arising from hiring experience, so we are studying managers who were willing to hire and manage black workers at least a few times over our sample period, which may represent employers with lower prior bias.

Lastly, we purposefully focus on a manager's own endogenous hiring decisions, since

[19] Along with other evidence presented above, these results are also inconsistent with managers learning about their own ability or evolving managing/screening ability over time as alternative explanations for our findings.

they should be more salient and managers should have a clear expectation of the worker from being in charge of hiring them. Hiring is fundamentally a selective process, and we want to understand how it is influenced by previous outcomes. We may worry that the effects we document could then be driven by worker and manager sorting. Yet, much of the previous reasoning applied against other alternatives applies to workers selectively applying for positions with managers based on their hiring record. In addition, Appendix C shows that negative experiences with black workers inherited, rather than hired, by the manager also decrease subsequent black hiring. More fundamentally, as mentioned in the setting section, workers apply for a job at the store or area level, typically do not know their manager until the interview, and are unlikely to observe information about the manager's hiring record until they are employed. Our results regarding the impact of other managers' experiences within a store, along with institutional details gathered through manager interviews, also suggest that workers are not assigned to a department manager based on the manager's previous experiences with worker groups, especially since we exclude transferred workers.

To summarize, the explanation which best jointly rationalizes our results is managers learning about the performance of black workers through their own hiring experiences with these workers. This learning could be quite broad, potentially including subjective productivity components and match quality. Still, the key takeaway is that managers aim to repeat experiences perceived as successful and avoid those perceived as unsuccessful. When they attribute some of the discrepancy between a worker's expected and realized productivity to potential differences between worker groups, then our theoretical framework predicts the creation and persistence of biased beliefs which generate the hiring patterns that we document.

7 Conclusion

We study the determinants of individual manager heterogeneity in the hiring of black workers using the employment records of a large US retailer and studying repeated hiring decisions of managers. We find that the hiring context, such as the location of a store, is an important determinant of the hiring of black workers across managers, but so are manager fixed effects and the previous hiring experiences of a manager with black workers.

We present a theoretical framework with the key proposition that managers learn and

update their beliefs about the expected performance of black workers from their experience on the job, and their updated beliefs affect their future hiring. We find empirical support for the main predictions of our framework. First, a manager's positive and negative experiences hiring black workers are correlated with future hiring from the group in the expected direction, suggesting that managers update their beliefs on the job. Second, managers react more following hiring experiences with black than white workers, consistent with them having less initial information about their performance and therefore weaker priors. Third, early negative experiences with black workers, have persistent impacts on subsequent hiring. This is consistent with negative experiences leading to negatively-biased beliefs which endogenously persist through decreased hiring and therefore decreased learning. Fourth, minority workers are systematically disadvantaged, because of differential updating by managers combined with the relative persistence of negative experiences on their subsequent hiring.

Our results have a number of implications for the literature, including for the study of discrimination and the organization of the firm. The fundamental issue is one of statistical inference from a selected sample of observations, leading managers to develop biased beliefs from their own experiences. Essentially, discrimination arises from a lack of learning or information.

Since inefficient hiring is costly for firms, this suggests scope for the design of organizational policies. First, policies could directly alter the hiring process to decrease discretion, for example by centralizing hiring at the store level and involving more than one manager in the process. Second, in theory, policies could also increase reliance on information about individual workers rather than their group(s), for example through pre-employment testing or trial periods, although the firm we study already has these in place and residual bias may still operate as long as individual performance is not fully revealed at hiring. Third, information could be aggregated and shared across managers so that they update from a larger set of experiences than their own, for example through hiring algorithms or store-level performance indicators. Fourth, more interaction could be induced between managers and minority workers, for example through affirmative action, cluster hiring, or internships. These may be particularly salient to managers since they directly lead them to interact with minority workers and learn about their performance, consistent with the contact hypothesis (Paluck et al., 2019). Our findings highlight the need for additional

evidence on the relationship between discrimination and organizational policies.

More fundamentally, the firm we study is an important employer, yet the firm's organization and the labor market in general appear to provide little corrective information to managers with individual idiosyncrasies in their hiring of minorities fueled by their personal experience. We present novel evidence of hiring discrimination attributable to endogenous employer learning about worker groups in a broadly representative labor market setting. Biased beliefs appear unlikely to resolve themselves through normal market interactions and at a minimum amount to several years of worse employment opportunities for minority workers. Classical models and tests of labor market discrimination are typically inadequate to capture discriminatory behavior arising from incorrect group perceptions, suggesting the importance of developing theories and gathering evidence on this more nuanced type of discrimination.

References

Abowd, John M, Francis Kramarz, and David N Margolis (1999) "High wage workers and high wage firms," *Econometrica*, 67 (2), 251–333.

Agarwal, Sumit, John C Driscoll, Xavier Gabaix, and David Laibson (2008) "Learning in the credit card market,"Technical report, National Bureau of Economic Research.

Aigner, Dennis J and Glen G Cain (1977) "Statistical theories of discrimination in labor markets," *Ilr Review*, 30 (2), 175–187.

Allport, Gordon Willard, Kenneth Clark, and Thomas Pettigrew (1954) *The nature of prejudice*: Addison-wesley Reading, MA.

Altonji, Joseph G and Charles R Pierret (2001) "Employer learning and statistical discrimination," *The quarterly journal of economics*, 116 (1), 313–350.

Arcidiacono, Peter, Patrick Bayer, and Aurel Hizmo (2010) "Beyond signaling and human capital: Education and the revelation of ability," *American Economic Journal: Applied Economics*, 2 (4), 76–104.

Arnold, David, Will Dobbie, and Crystal S Yang (2018) "Racial bias in bail decisions," *The Quarterly Journal of Economics*, 133 (4), 1885–1932.

Arrow, Kenneth (1973) "The theory of discrimination," *Discrimination in Labor Markets*, 3 (10), 3–33.

Åslund, Olof, Lena Hensvik, and Oskar Nordström Skans (2014) "Seeking similarity: How immigrants and natives manage in the labor market," *Journal of Labor Economics*, 32 (3), 405–441.

Autor, David H and David Scarborough (2008) "Does job testing harm minority workers? Evidence from retail establishments," *The Quarterly Journal of Economics*, 123 (1), 219–277.

Becker, Gary S (1957) *The Economics of Discrimination*: University of Chicago press.

Benson, Alan, Simon Board, and Moritz Meyer-ter Vehn (2019) "Discrimination in hiring: Evidence from retail sales," Unpublished.

Bergman, Jukka-Pekka, Pasi Luukka, Ari Jantunen, and Anssi Tarkiainen (2020) "Cognitive diversity, managerial characteristics and performance differences across the cleantech firms," *International Journal of Knowledge-Based Organizations (IJKBO)*, 10 (1), 1–26.

Berson, Clémence, Morgane Laouenan, and Emmanuel Valat (2020) "Outsourcing recruitment as a solution to prevent discrimination: A correspondence study," *Labour Economics*, 64, 101838.

Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan (2005) "Implicit discrimination," *American Economic Review*, 95 (2), 94–98.

Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope (2021) "Inaccurate statistical discrimination,"Technical Report w25935, National Bureau of Economic Research.

Bohren, J Aislinn, Alex Imas, and Michael Rosenberg (2019) "The dynamics of discrimination: Theory and evidence," *American economic review*, 109 (10), 3395–3436.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2016) "Stereotypes," *The Quarterly Journal of Economics*, 131 (4), 1753–1794.

——— (2019) "Beliefs about gender," *American Economic Review*, 109 (3), 739–73.

Boushey, Heather and Sarah Jane Glynn (2012) "There are significant business costs to replacing employees," *Center for American Progress*, 16, 1–9.

Bureau of Labor Statistics (2021) "Retail Trade: NAICS 44-45. Industries at a Glance. Available at https://www.bls.gov/iag/tgs/iag44-45.htm [Date Published: 02/10/2021] [Date Accessed: 02/11/2021]."

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Housman (2015) "The value of hiring through employee referrals," *The Quarterly Journal of Economics*, 130 (2), 805–839.

Bycio, Peter, Rick D Hackett, and Kenneth M Alvares (1990) "Job performance and turnover: a review and meta-analysis," *Applied Psychology*, 39 (1), 47–76.

Charles, Kerwin Kofi and Jonathan Guryan (2011) "Studying discrimination: Fundamental challenges and recent progress," *Annu. Rev. Econ.*, 3 (1), 479–511.

Coate, Stephen and Glenn C Loury (1993) "Will affirmative-action policies eliminate negative stereotypes?" *The American Economic Review*, 1220–1240.

Cornell, Bradford and Ivo Welch (1996) "Culture, information, and screening discrimination," *Journal of political Economy*, 104 (3), 542–571.

Cullen, Zoë and Ricardo Perez-Truglia (2021) "The Old Boys' Club: Schmoozing and the Gender Gap,"Technical Report w24841, National Bureau of Economic Research.

Dipboye, Robert L, Kevin Wooten, and Stefanie K Halverson (2004) "Behavioral and situational interviews," in Thomas, Jay C ed. *Comprehensive handbook of psychological assessment*, 4, 297–316, Hoboken, NJ, US: John Wiley & Sons, Inc.

Erev, Ido and Ernan Haruvy (2016) "Learning and the Economics of Small Decisions," in *The Handbook of Experimental Economics, Volume 2*, 638–716: Princeton University Press.

Farber, Henry S and Robert Gibbons (1996) "Learning and wage dynamics," *The Quarterly Journal of Economics*, 111 (4), 1007–1047.

Gallagher, Justin (2014) "Learning about an infrequent event: evidence from flood insurance take-up in the United States," *American Economic Journal: Applied Economics*, 206–233.

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.

Giuliano, Laura and Michael R Ransom (2013) "Manager ethnicity and employment segregation," *ILR Review*, 66 (2), 346–379.

Glover, Dylan, Amanda Pallais, and William Pariente (2017) "Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores," *The Quarterly Journal of Economics*, 132 (3), 1219–1260.

Guarino, Cassandra, Mark Reckase, Brian Stacy, and Jeffrey Wooldridge (2015) "A comparison of student growth percentile and value-added models of teacher performance," *Statistics and Public Policy*, 2 (1), 1–11.

Guryan, Jonathan and Kerwin Kofi Charles (2013) "Taste-based or statistical discrimination: the economics of discrimination returns to its roots," *The Economic Journal*, 123 (572), F417–F432.

Hjort, Jonas (2014) "Ethnic divisions and production in firms," *The Quarterly Journal of Economics*, 129 (4), 1899–1946.

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li (2018) "Discretion in hiring," *The Quarterly Journal of Economics*, 133 (2), 765–800.

Huang, Chung, Hao-Chieh Lin, and Chih-Hsun Chuang (2006) "Constructing factors related to worker retention," *International Journal of Manpower*, 27 (5), 491–508.

Hull, Peter (2021) "What Marginal Outcome Tests Can Tell Us About Racially Biased Decision-Making,"Technical Report w28503, National Bureau of Economic Research.

Kahn, Lisa B and Fabian Lange (2014) "Employer learning, productivity, and the earnings distribution: Evidence from performance measures," *The Review of Economic Studies*, 81 (4), 1575–1613.

Kirnan, Jean Powell, John A Farley, and Kurt F Geisinger (1989) "The relationship between recruiting source, applicant quality, and hire performance: An analysis by sex, ethnicity, and age," *Personnel psychology*, 42 (2), 293–308.

Lang, Kevin (1986) "A language theory of discrimination," *The quarterly journal of economics*, 101 (2), 363–382.

Lang, Kevin and Jee-Yeon K Lehmann (2012) "Racial discrimination in the labor market: Theory and empirics," *Journal of Economic Literature*, 50 (4), 959–1006.

Lange, Fabian (2007) "The speed of employer learning," *Journal of Labor Economics*, 25 (1), 1–35.

Lazear, Edward P, Kathryn L Shaw, and Christopher Stanton (2016) "Making do with less: working harder during recessions," *Journal of Labor Economics*, 34 (S1), S333–S360.

Lepage, Louis Pierre (2021a) "Endogenous learning, persistent employer biases, and discrimination," Unpublished.

——— (2021b) "Stereotype Formation and Hiring Discrimination," Unpublished.

Leung, Ming D (2018) "Learning to hire? Hiring as a dynamic experiential learning process in an online market for contract labor," *Management Science*, 64 (12), 5651–5668.

Lundberg, Shelly J and Richard Startz (1983) "Private discrimination and social intervention in competitive labor market," *The American Economic Review*, 73 (3), 340–347.

Malmendier, Ulrike (2021a) "Experience effects in finance: Foundations, applications, and future directions," *Review of Finance*, 25 (5), 1339–1363.

——— (2021b) "Exposure, experience, and expertise: Why personal histories matter in economics," *Journal of the European Economic Association*.

Malmendier, Ulrike and Stefan Nagel (2011) "Depression babies: do macroeconomic experiences affect risk taking?" *The quarterly journal of economics*, 126 (1), 373–416.

Montgomery, James D (1991) "Social networks and labor-market outcomes: Toward an economic analysis," *The American economic review*, 81 (5), 1408–1418.

Morgan, John and Felix Várdy (2009) "Diversity in the Workplace," *American Economic Review*, 99 (1), 472–85.

Morris, Carl N (1983) "Parametric empirical Bayes inference: theory and applications," *Journal of the American statistical Association*, 78 (381), 47–55.

Paluck, Elizabeth Levy, Seth A Green, and Donald P Green (2019) "The contact hypothesis re-evaluated," *Behavioural Public Policy*, 3 (2), 129–158.

Phelps, Edmund S (1972) "The statistical theory of racism and sexism," *The american economic review*, 62 (4), 659–661.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales (2014) "How stereotypes impair women's careers in science," *Proceedings of the National Academy of Sciences*, 111 (12), 4403–4408.

Ronchi, Maddalena and Nina Smith (2021) "Daddy's girl: Daughters, managerial decisions, and gender inequality," Unpublished.

Rooth, Dan-Olof (2010) "Automatic associations and discrimination in hiring: Real world evidence," *Labour Economics*, 17 (3), 523–534.

Sarsons, Heather (2019) "Interpreting signals in the labor market: evidence from medical referrals," Unpublished.

Stephens-Davidowitz, Seth (2014) "The cost of racial animus on a black candidate: Evidence using Google search data," *Journal of Public Economics*, 118, 26–40.

Sue, Derald Wing (2006) *The Invisible Whiteness of Being: Whiteness, White Supremacy, White Privilege, and Racism.*: John Wiley & Sons Inc.

Taber, Mary E and Wallace Hendricks (2003) "The effect of workplace gender and race demographic composition on hiring through employee referrals," *Human Resource Development Quarterly*, 14 (3), 303–319.

Williams, Charles R and Linda Parrack Livingstone (1994) "Another look at the relationship between performance and voluntary turnover," *Academy of Management Journal*, 37 (2), 269–298.

Zimmerman, Ryan D and Todd C Darnold (2009) "The impact of job performance on employee turnover intentions and the voluntary turnover process: A meta-analysis and path model," *Personnel review*, 38 (2), 142–158.

Figure 1: Cumulative turnover by tenure

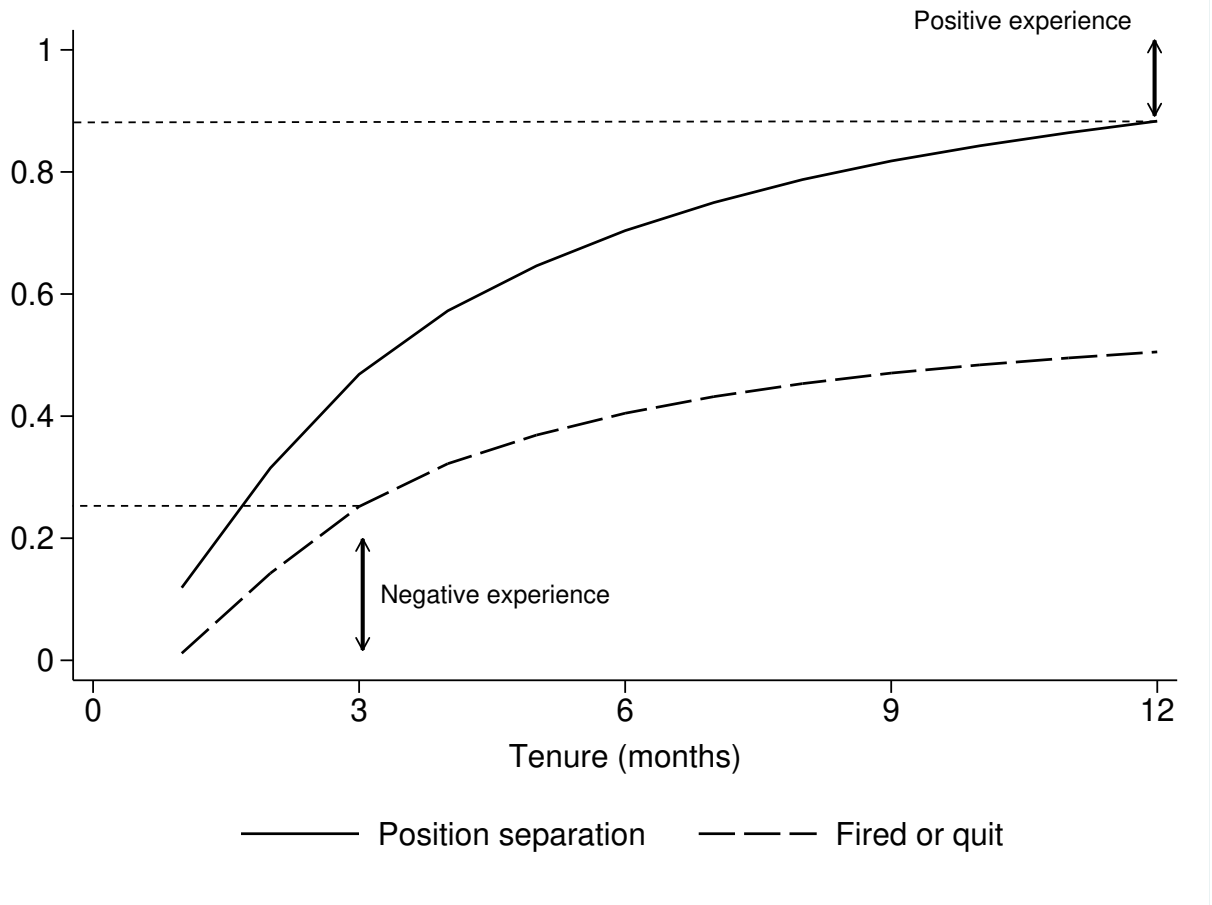

NOTE. "Position separation" refers to the worker no longer working in the position for which they were hired, including dismissals and quits but also department or store transfers, layoffs, promotions, and retirement/disability/death. The sample is restricted to workers hired at least one year before the end of our sample period.

Figure 2: Kernel density estimates of manager hiring shares, black and white workers

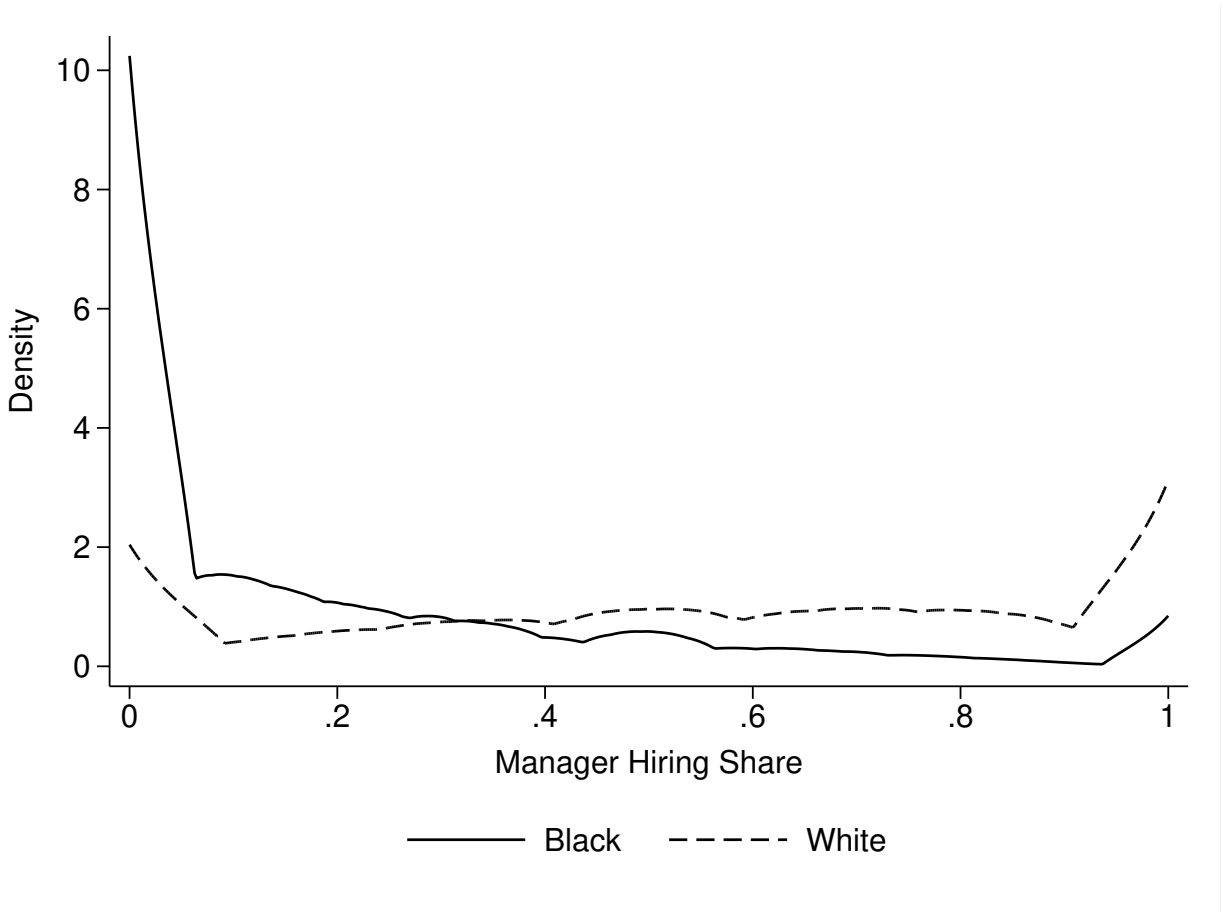

NOTE. Black share bandwidth: 0.066. White share bandwidth: 0.091.

Figure 3: Impact of a manager's first hiring experience with black and white workers on their subsequent black hiring share

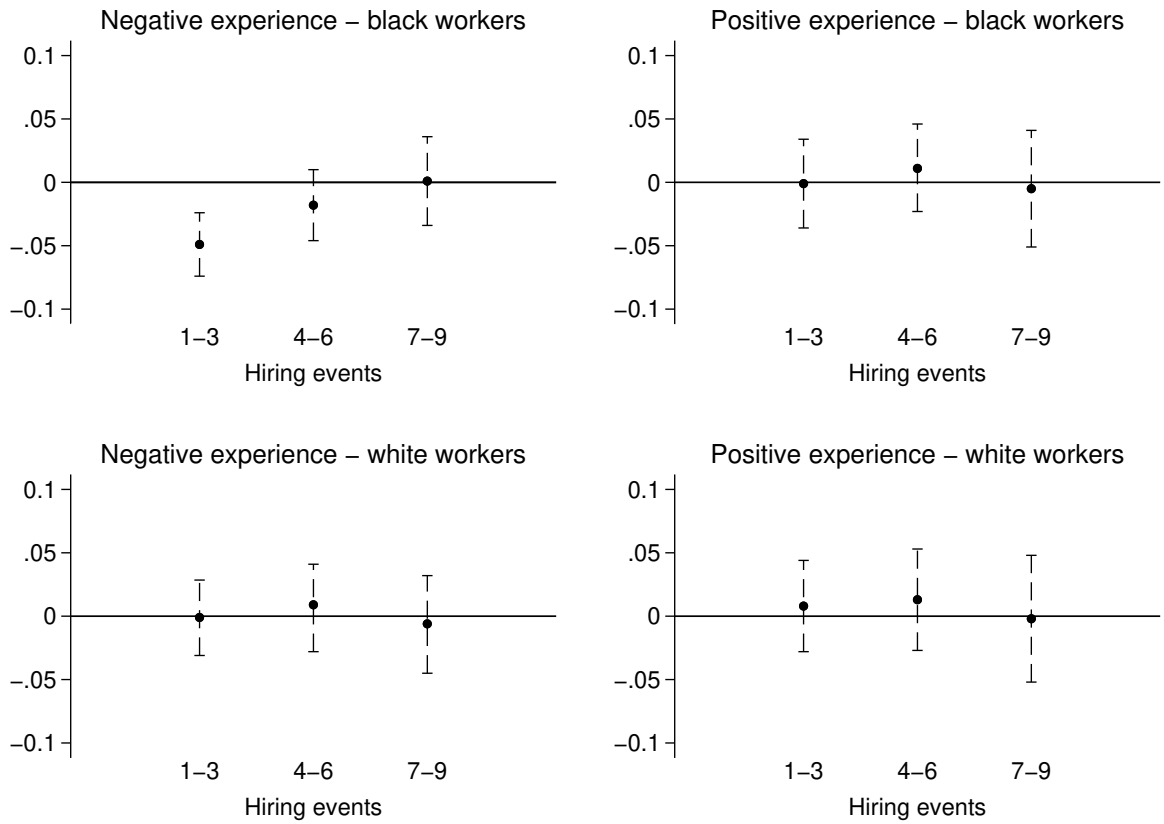

NOTE. 95% confidence intervals from clustered standard errors at the manager level are presented using dashed lines. A negative experience corresponds to the share of a manager's first hire(s) from a racial group that was fired or quit in the first 3 months of employment. A positive experience corresponds to the share of a manager's first hire(s) from a racial group that achieved tenure of at least one year in their position. Regressions include the fraction of full-time and female hires, average age of hires, total number of workers hired in the event, yearly unemployment and college attainment rates in the state, month and year, and store fixed effects.

Table 1: Summary statistics and performance measures

	Workers		Managers	
	Mean	Std. dev.	Mean	Std. dev.
Age	31.01	(14.14)	41.82	(11.39)
Female	0.56	(0.50)	0.37	(0.48)
White	0.55	(0.50)	0.73	(0.44)
Black	0.22	(0.41)	0.11	(0.31)
Tenure	33.73	(71.68)	122.75	(125.17)
Full time	0.18	(0.37)	0.99	(0.09)
N. hires			30.16	(65.75)
N. persons	1,067,682		27,470	
N. person-months	17,445,003		684,218	
Fired or quit within 3 months (Black)	0.270	(0.444)		
Fired or quit within 3 months (White)	0.250	(0.432)		
Tenure above 1 year (Black)	0.116	(0.321)		
Tenure above 1 year (White)	0.145	(0.352)		

NOTE. Performance measures are calculated at the individual hire level. Tenure corresponds to tenure in the position for which the worker was hired. The absence of a worker quitting or being fired within one year does not imply that the worker has achieved a year of tenure, given transfers and layoffs.

Table 2: OLS estimates of the cumulative impact of previous experiences with black and white workers on current black hiring, expected tenure

Black fraction hired	Restricted to managers observed at tenure = 1 (1)	Full panel (2)
Black expected months of tenure	0.050	0.060
	(0.011)	(0.008)
White expected months of tenure	-0.009	-0.012
	(0.011)	(0.007)
Manager FE	Y	Y
Store FE	Y	Y
Hiring month and year FE	Y	Y
Worker and event controls	Y	Y
Outcome mean	0.380	0.365
P-value: B = -1 * W	0.023	0.000
Standard deviation black	0.524	0.496
Standard deviation white	0.555	0.527
Observations	35,937	72,424

NOTE. Clustered standard errors at the manager level are presented in parentheses. One observation corresponds to a manager-month in which at least one worker was hired. Expected months of tenure corresponds to the cumulative average deviation from expected tenure at the firm for workers hired by the manager. Regressions include the fraction of full-time and female hires, average age of hires, total number of workers hired in the event, number of previous hiring events, time since the last hiring event, yearly unemployment and college attainment rates in the state, month and year, store, and manager fixed effects. Column 1 restricts the sample to managers who started hiring at the firm for the first time during our sample period.

Table 3: OLS estimates of the cumulative impact of previous experiences with black and white workers on current black hiring, negative and positive experiences

	Manager observed at tenure = 1	Full panel	Manager observed at tenure = 1	Full panel
Black fraction hired	(1)	(2)	(3)	(4)
Black fraction quit/fired $\leq$ 3 months	-0.072	-0.067		
	(0.017)	(0.011)		
White fraction quit/fired $\leq$ 3 months	0.044	0.067		
	(0.022)	(0.016)		
Black fraction tenure $\geq$ 12 months			0.058	0.062
			(0.024)	(0.016)
White fraction tenure $\geq$ 12 months			-0.001	-0.015
			(0.027)	(0.018)
Manager FE	Y	Y	Y	Y
Store FE	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y
Outcome mean	0.348	0.337	0.347	0.336
P-value: B = -1 * W	0.077	0.169	0.090	0.025
Standard deviation black	0.294	0.282	0.229	0.230
Standard deviation white	0.207	0.192	0.205	0.204
Observations	31,911	66,692	26,655	56,911

NOTE. Clustered standard errors at the manager level are presented in parentheses. See Table 2 for additional notes.

Table 4: OLS estimates of the persistence of previous negative experiences with black and white workers on current black hiring

	Manager observed at tenure = 1			Full panel		
Black fraction hired	Event 1 (1)	Events 1-3 (2)	Event t-1 (3)	Event 1 (4)	Events 1-3 (5)	Event t-1 (6)
Black fraction quit/fired $\leq$ 3 months	-0.019	-0.042	-0.041	-0.007	-0.024	-0.043
	(0.007)	(0.014)	(0.006)	(0.004)	(0.009)	(0.004)
White fraction quit/fired $\leq$ 3 months	-0.009	-0.013	0.035	0.000	-0.008	0.036
	(0.008)	(0.017)	(0.006)	(0.005)	(0.011)	(0.004)
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y	Y	Y
Manager FE			Y			Y
Outcome mean	0.367	0.402	0.348	0.354	0.385	0.337
P-value: B = -1 * W	0.006	0.005	0.174	0.198	0.010	0.020
Standard deviation black	0.405	0.239	0.409	0.407	0.238	0.408
Standard deviation white	0.348	0.205	0.352	0.341	0.200	0.352
Observations	35,613	27,829	31,911	72,880	59,762	66,692

NOTE. Clustered standard errors at the manager level are presented in parentheses for columns 3 and 6, robust standard errors are presented in parentheses for columns 1-2 and 4-5. Event 1 refers to the first group of workers from a given race hired by the manager, Events 1-3 refers to the average outcome of the first three groups (excluding managers who hired less than three groups of workers from either racial group). These specifications do not include manager fixed effects since the nature of a manager's first experience(s) is time invariant. Event t-1 refers to the latest group of hires from a given race and these specifications also include manager fixed effects. See Table 2 for additional details.

Table 5: OLS estimates of the persistence of previous positive experiences with black and white workers on current black hiring

	Manager observed at tenure = 1			Full panel		
Black fraction hired	Event 1 (1)	Events 1-3 (2)	Event t-1 (3)	Event 1 (4)	Events 1-3 (5)	Event t-1 (6)
Black fraction tenure $\geq$ 12 months	-0.010	0.020	0.022	-0.012	0.006	0.019
	(0.011)	(0.018)	(0.009)	(0.006)	(0.012)	(0.006)
White fraction tenure $\geq$ 12 months	-0.003	0.018	-0.016	0.002	0.002	-0.006
	(0.011)	(0.020)	(0.008)	(0.006)	(0.012)	(0.005)
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y	Y	Y
Manager FE			Y			Y
Outcome mean	0.366	0.401	0.347	0.353	0.385	0.336
P-value: B = -1 * W	0.375	0.146	0.658	0.271	0.645	0.080
Standard deviation black	0.327	0.200	0.296	0.339	0.206	0.306
Standard deviation white	0.342	0.214	0.297	0.354	0.224	0.309
Observations	29,869	23,527	27,249	62,487	51,696	58,026

NOTE. Clustered standard errors at the manager level are presented in parentheses for columns 3 and 6, robust standard errors are presented in parentheses for columns 1-2 and 4-5. Event 1 refers to the first group of workers from a given race hired by the manager, Events 1-3 refers to the average outcome of the first three groups (excluding managers who hired less than three groups of workers from either racial group). These specifications do not include manager fixed effects since the nature of a manager's first experience(s) is time invariant. Event t-1 refers to the latest group of hires from a given race and these specifications also include manager fixed effects. See Table 2 for additional details.

Table 6: OLS estimates of the cumulative impact of negative experiences on black hiring throughout a manager's hiring history

Black fraction hired	Early experiences (1)	Middle experiences (2)	Late experiences (3)
Black fraction quit/fired $\leq$ 3 months	-0.096	-0.082	-0.068
	(0.037)	(0.036)	(0.036)
Manager FE	Y	Y	Y
Store FE	Y	Y	Y
Hiring month and year FE	Y	Y	Y
Worker and event controls	Y	Y	Y
Outcome mean	0.382	0.419	0.428
Standard deviation	0.335	0.324	0.313
Observations	6,999	7,347	6,272

NOTE. Each manager's previous hiring events are separated chronologically into three experience terciles. See Table 2 for additional details.

Table 7: OLS estimates of the serial correlation of negative and positive experiences with black workers

	Fraction quit/fired $\leq$ 3 months (1)	Fraction tenure $\geq$ 12 months (2)
Lagged fraction quit/fired $\leq$ 3 months	-0.343	
	(0.032)	
Lagged fraction tenure $\geq$ 12 months		-0.369
		(0.036)
Manager FE	Y	Y
Store FE	Y	Y
Hiring month and year FE	Y	Y
Worker and event controls	Y	Y
Outcome mean	0.231	0.111
Standard deviation	0.227	0.202
Observations	13,222	11,208

NOTE. See Table 2 for additional details.

Appendix A Worker performance and tenure

In the absence of direct measures of positive and negative signals observed by managers about their hires, we use measures indicating higher or lower-than-expected turnover at the firm. While managers also value minimizing turnover rather than simply maximizing the productivity of their hires, we investigate the extent to which tenure captures worker performance to some extent, providing a valid proxy for productivity for individual productivity.

A direct measure of individual productivity could be valuable, in case a highly productive hire who stays for a shorter time may be preferable to a worse hire who stays longer. Unfortunately, relatively few workers in our data have objective individual productivity measures, and for those who do, they can be inconsistent. Given our sample restrictions, the only direct performance measure applies to 61,665 floor salespeople who are evaluated on the revenue of their sales versus targets.[20]

As shown in Table A, sales performance statistically significantly increases with tenure, even given the fact that sales targets also increase with tenure, reflecting that workers who achieve longer tenure are on average more productive (approximately 5% more per year of tenure). Similarly, the likelihood of a worker being promoted to a manager in a given month unsurprisingly increases with tenure, showing a positive relationship between tenure and match quality. Moreover, involuntary terminations also decrease with tenure, in particular those arising from unsatisfactory performance and job abandonment without warning, which are particularly costly to the firm. In contrast, the probability of a termination being voluntary, specifically for career advancement or return to studies, increases with tenure, suggesting that workers who stay longer are more productive.

Overall, across several measures of performance and productivity, these results support our analysis using shorter or longer tenure as the basis for negative and positive productivity signals.

[20] Although there are numerically many salespeople, the unit of observation in the analysis is their superior who forms beliefs before making hiring decisions, leading to power and dimensionality problems. We also investigated other measures of positive experiences including salary and additional work hours. Entry salaries are generally fixed or pre-determined and additional hours primarily reflect demand fluctuations.

Table A: Worker performance and tenure

	Sales performance (1)	Promotion to manager (2)	Involuntary termination (3)	Termination for performance or abandonment (4)	Termination for career or studies (5)
Months of tenure	70.832	0.00007	-0.0034	-0.0165	0.0159
	(7.042)	(0.00001)	(0.0003)	(0.0004)	(0.0003)
Manager FE	Y		Y	Y	Y
Store FE	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y
Job title FE	Y	Y	Y	Y	Y
Outcome mean	17,596	0.0026	0.2000	0.2334	0.3302
Observations	122,967	626,544	67,212	67,212	67,212

NOTE. Robust standard errors are presented in parentheses. Sales performance is the total monthly commissioned sales in dollars relative to a target. Promotion to manager is the probability that a worker be promoted to manager in a given month. Termination for performance or abandonment includes involuntary terminations specifically for unsatisfactory worker performance or because the worker abandons their job. Termination for career or studies includes voluntary terminations specifically for the worker's career advancement or return to studies.

Appendix B Variation in black hiring across managers

We describe heterogeneity in the hiring of black workers across managers and examine how much of it is due to idiosyncratic variation across individual managers versus external factors. Theories of discrimination fundamentally differ along this dimension. Under classical statistical discrimination, managers discriminate similarly around the true productivity distribution of each group; they are not idiosyncratically biased. In stark contrast, idiosyncratic prejudice or bias are at the center of taste-based and belief-based discrimination.

Many factors presumably contribute to this heterogeneity, such as store location. To estimate how much heterogeneity in black hiring is explained by manager effects net of other factors that may vary by store, department, job, time period, or economic condition, we take Abowd et al. (1999)'s approach of analyzing connected sets of workers.[21] Over a quarter of managers hire in more than one store, around 8% hire in more than 2 stores, and the majority of managers hire for multiple job types, generating substantial variation to separately identify manager fixed effects. Indeed, the largest connected set of managers and stores covers over 90% of new workers hired at the firm during our sample period.

We implement this approach using a linear probability model of the form

$$Black_{imjlt} = X_{mjlt}\beta + \gamma_m + \alpha_j + \lambda_l + \theta_t + \varepsilon_{imjlt} \tag{3}$$

where the dependent variable indicates that worker i hired by manager m for job j in location l at time t is black. X_{mjlt} includes whether the worker was hired for a part-time or full-time job, the manager's cumulative number of hires, the yearly state unemployment rate, and the fraction of the state population with at least some college education. γ_m, α_j, λ_l, and θ_t correspond to manager, job, store, and month and year fixed effects.[22] We compute the predicted value for each individual hire and average predicted values at the manager level to obtain the predicted share of black hires for each manager. This procedure yields higher predicted shares for managers recruiting in jobs, locations, periods, and market conditions associated with more black hires.

[21] Several recent papers have also applied this approach to estimate manager fixed effects net of sets of highly correlated covariates e.g. Lazear et al. (2016) and Benson et al. (2019)

[22] The results are similar when including department fixed effects as well as worker demographics including age and gender.

Figure B contrasts the predicted black hiring shares across managers with the actual values. By construction, predicted shares approximate the middle of the distribution. Especially without manager fixed effects, they fail to capture much of the bottom of the distribution, predicting that too many managers hire 10-30% black workers and too few hire less. Beyond manager fixed effects, the majority of the explanatory power comes from the store fixed effects, which capture store and area-level characteristics. Manager fixed effects alone explain 4-5% of the total variation in black hiring and roughly a third of the discrepancy between actual shares and those predicted by the model without manager fixed effects. Qualitatively, the model with manager fixed effects still under-predicts the share of managers who hire very few or no black workers, but the discrepancy is substantively smaller. This exercise suggests that, beyond store and contextual factors, the specific identify of the hiring manager is an important predictor of black hiring in a department. Figure B presents analogous results for white hiring while Figure B presents results restricted to managers who hire at least 5 workers over our sample period, highlighting that manager fixed effects explain a particularly large share of residual variation in black hiring for that subset of managers.

The distribution of manager fixed effects is shown in Figure B. To adjust the estimated fixed effects based on their precision from the total number of hires by each manager, we apply an empirical Bayes shrinkage procedure, although its impact on the estimates is negligible (Morris, 1983; Guarino et al., 2015). The distribution appears fairly symmetric. As shown in Figure **??**, the analogous distribution for white workers exhibits a slight positive skew. Simple correlation analyses indicate that the fixed effects for black hiring are negatively correlated with turnover of black workers, suggesting that they capture something concrete about the ability or willingness of managers to successfully hire and manage these workers. In contrast, there is little correlation between the fixed effects and the state-level prejudice measure from Stephens-Davidowitz (2014) after controlling for the fraction of black population in the Core-Based Statistical Area (CBSA).

Figure B1: Kernel density estimates of predicted black hiring shares

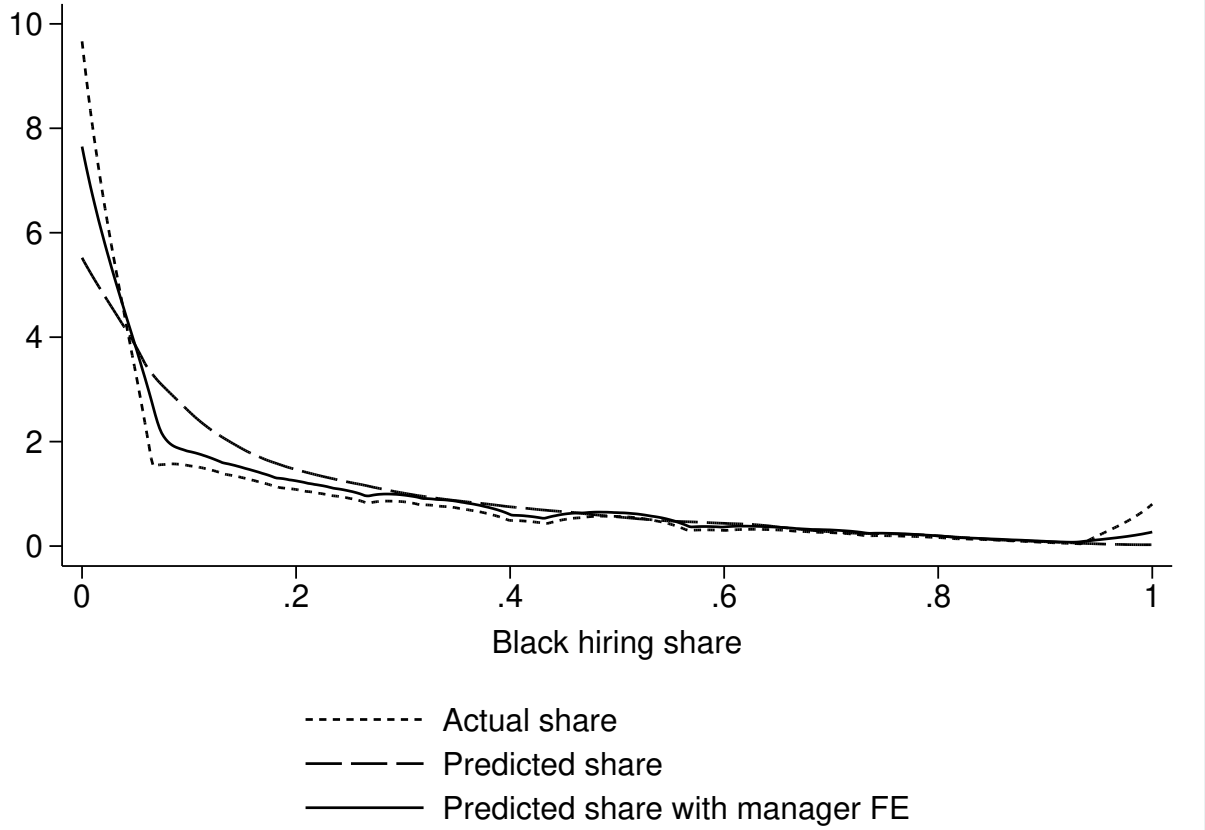

NOTE. Predicted shares are obtained by averaging predicted values for each manager from an individual hire level linear probability model regression including whether the worker was hired for a part-time or full-time job, the manager's previous number of hires at the time that the current worker is hired, yearly state unemployment rate and fraction with at least some college education, and month and year, store, job title, and individual manager fixed effects. A small fraction of predicted values outside of the 0-1 range were replaced with values of 0 or 1 for ease of visualization. Actual share bandwidth: 0.066. Predicted share bandwidth: 0.059. Predicted share with manager FE bandwidth: 0.068.

Figure B2: Kernel density estimates of manager predicted white hiring shares

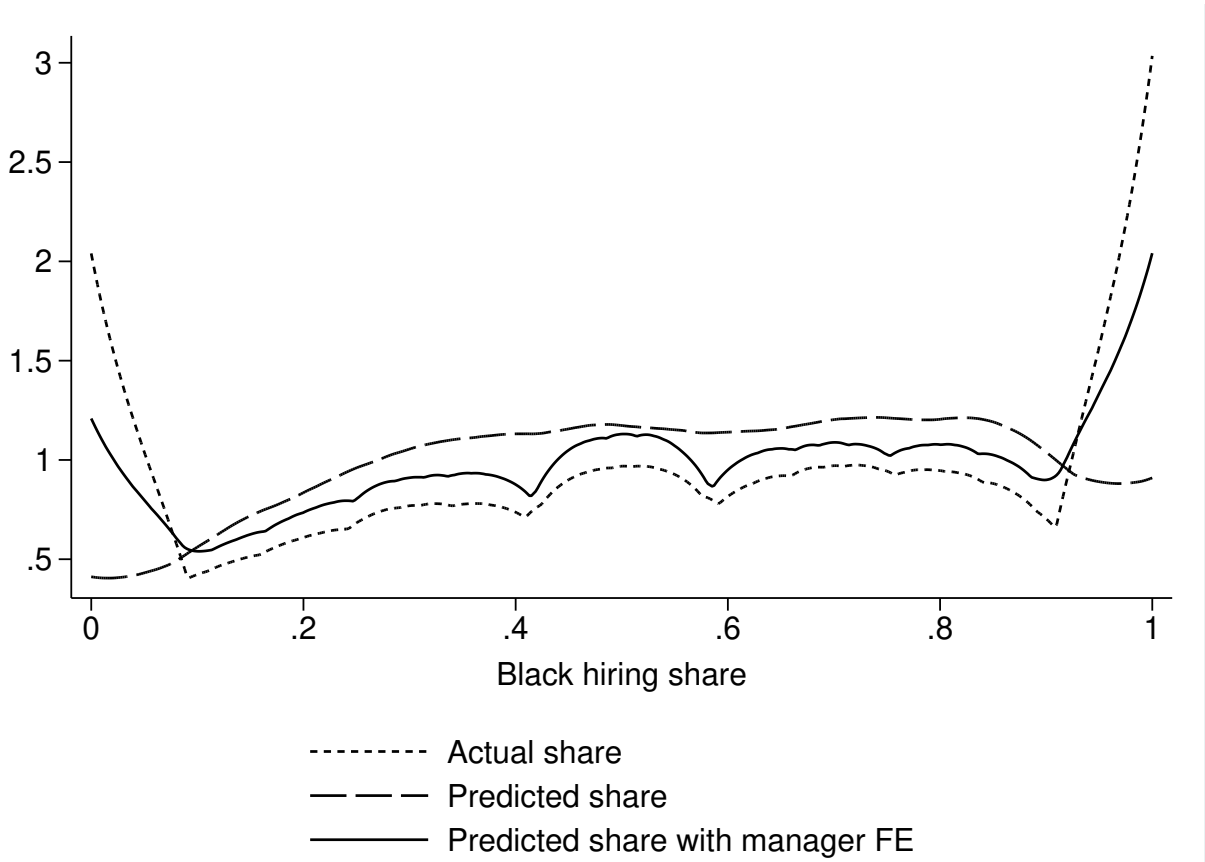

NOTE. See Figure 2 for details. Actual share bandwidth: 0.091. Predicted share bandwidth: 0.074. Predicted share with manager FE bandwidth: 0.086.

Figure B3: Kernel density estimates of manager predicted black hiring shares, managers with over 5 hires

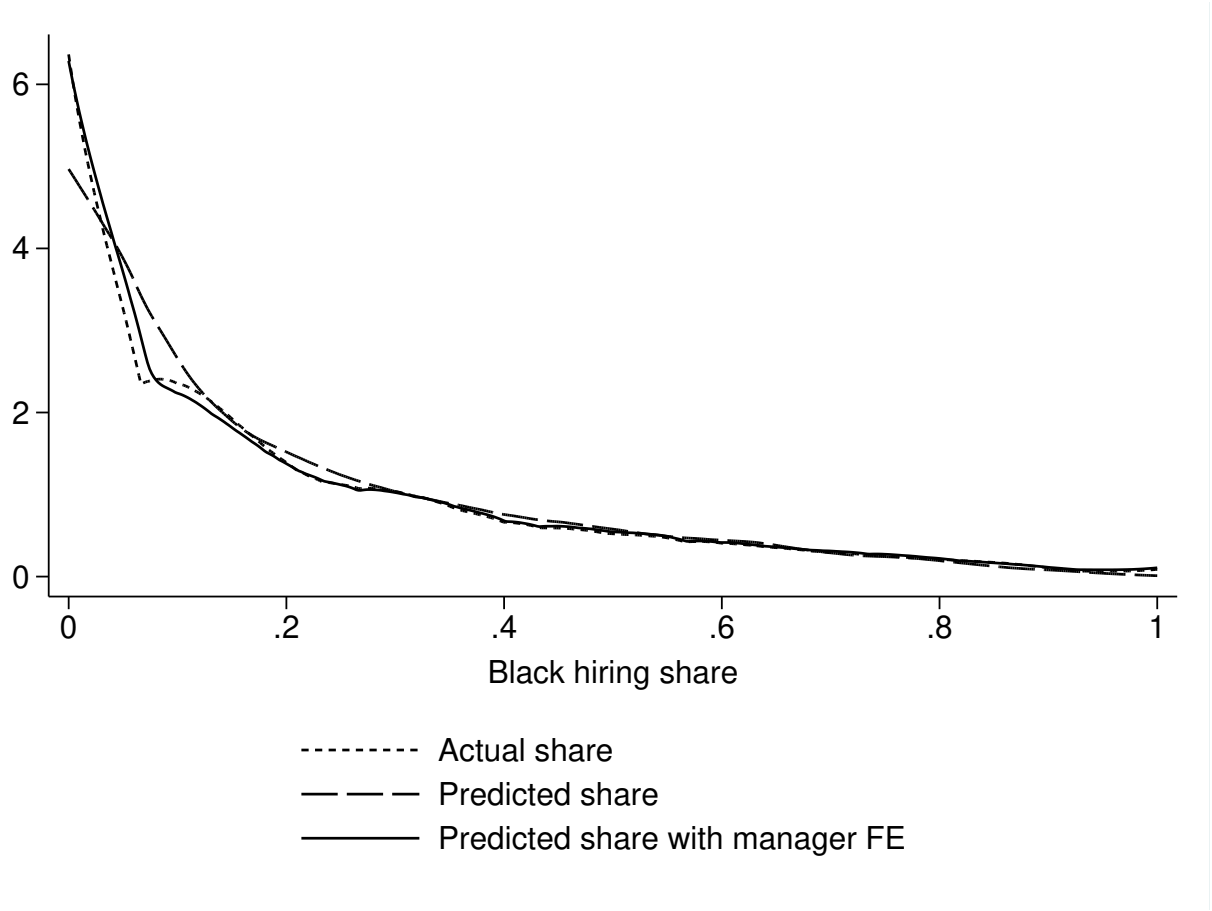

NOTE. See Figure 2 for details. Actual share bandwidth: 0.066. Predicted share bandwidth: 0.063. Predicted share with manager FE bandwidth: 0.068.

Figure B4: Distribution of manager fixed effects for black hiring

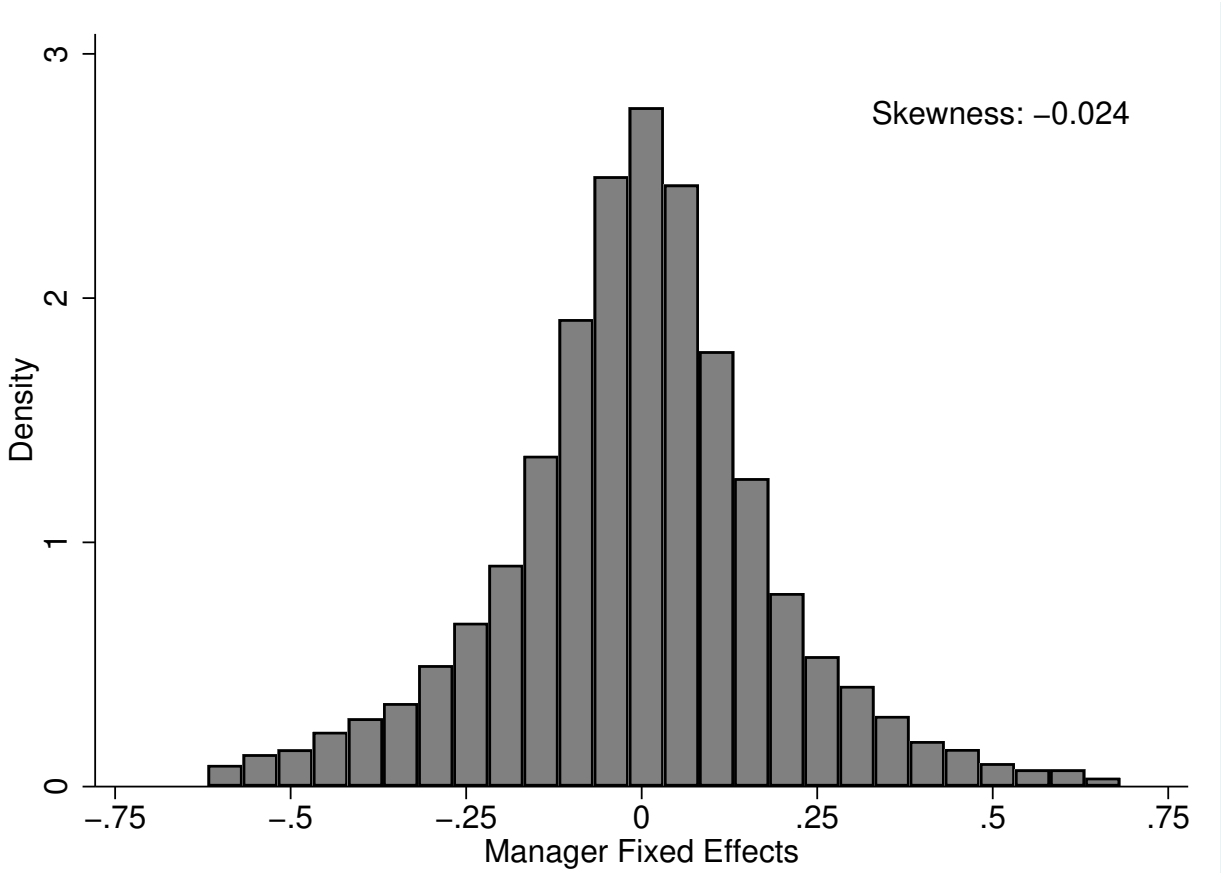

NOTE. See Figure 2 for specification details. Fixed effects are estimated for the largest connected sample of stores and managers following Abowd et al. (1999) and adjusted using empirical Bayes shrinkage.

Figure B5: Distribution of manager fixed effects for white hiring

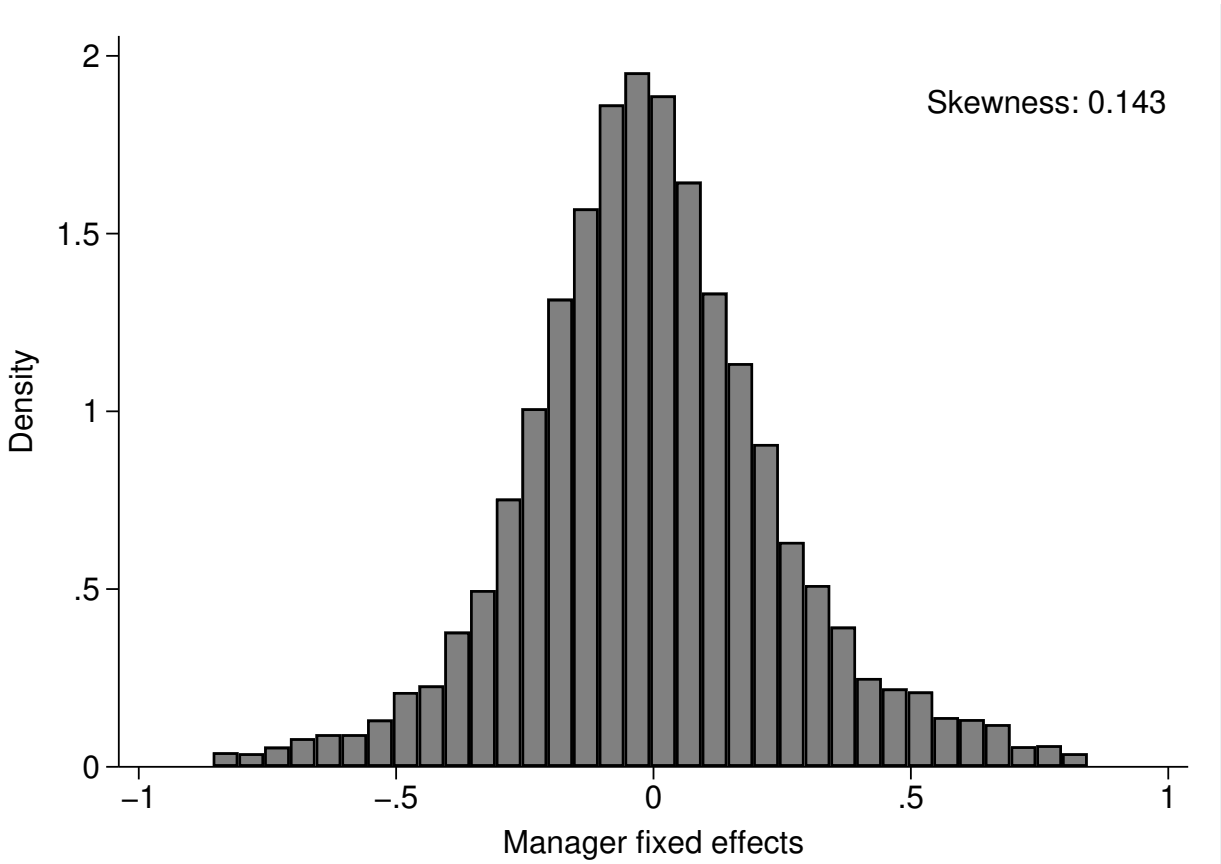

Note. See Figures 2 and 3 for details.

Appendix C Additional results on the impact of hiring experiences

Table C1: Additional Experience Measures

Black fraction hired	Fired (1)	Quit (2)	Relative to white (3)	Relative to CBSA (4)
Black fraction quit/fired $\leq$ 3 months	-0.087	-0.057	-0.068	-0.079
	(0.027)	(0.018)	(0.014)	(0.016)
Manager FE	Y	Y	Y	Y
Store FE	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y
Outcome mean	0.367	0.37	0.348	0.369
Observations	33,971	33,971	31,911	33,675

NOTE. Clustered standard errors at the manager level are presented in parentheses. See Table 2 for additional details.

Table C2: Sources of Updating

Black fraction hired	"Exogenous" separation (1)	"Endowed" workers (2)	Store change (3)	Avg. exp. versus others' (4)
Black fraction quit/fired $\leq$ 3 months	-0.063	-0.034	-0.067	-0.069
	(0.032)	(0.027)	(0.045)	(0.023)
Other managers at the store				
Black fraction quit/fired $\leq$ 3 months				-0.016
				(0.025)
White fraction quit/fired $\leq$ 3 months				-0.004
				(0.037)
Manager FE	Y	Y	Y	Y
Store FE	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y
Outcome mean	0.367	0.328	0.356	0.352
Observations	33,971	10,911	977	30,985

NOTE. Clustered standard errors at the manager level are presented in parentheses. "Exogenous separation" restricts fires and quits to dissatisfaction with pay, compensation or benefits, worker integrity, illegal or unethical behavior, or violation of rules and policies. "Endowed workers" corresponds to workers already in the department at the manager's arrival. See Table 2 for additional details.

Table C3: OLS estimates of the cumulative impact of previous experiences with black and white workers on current black hiring, bottom and top quartiles of tenure

	Manager observed at tenure = 1	Full panel	Manager observed at tenure = 1	Full panel
Black fraction hired	(1)	(2)	(3)	(4)
Black expected tenure in the bottom quartile	-0.050	-0.053		
	(0.007)	(0.005)		
White expected tenure in the bottom quartile	0.022	0.021		
	(0.007)	(0.005)		
Black expected tenure in the top quartile			0.065	0.062
			(0.007)	(0.005)
White expected tenure in the top quartile			0.026	-0.001
			(0.007)	(0.005)
Manager FE	Y	Y	Y	Y
Store FE	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y
Outcome mean	0.380	0.365	0.380	0.365
Observations	35,883	72,424	35,883	72,424

NOTE. Clustered standard errors at the manager level are presented in parentheses. Expected tenure corresponds to the cumulative average deviation from expected tenure at the firm for workers hired by the manager. See Table 2 for additional details.

Table C4: OLS estimates of a time placebo for the first hiring experience on black hiring

Black fraction hired	Placebo event 1 (1)
Black fraction quit/fired $\leq$ 3 months	0.012
	(0.047)
Manager FE	Y
Store FE	Y
Hiring month and year FE	Y
Worker and event controls	Y
Outcome mean	0.284
Observations	4,738

NOTE. Robust standard errors are presented in parentheses. The coefficient represents a time placebo investigating the impact of a worse performance by black hires in the month before a manager begins in their position at the department and the outcome is the hiring of black workers by the manager in subsequent hiring events. See Table 2 for additional details.

Table C5: OLS estimates of the correlation between the existing share of black workers in a department and the black hiring share

Black fraction hired	(1)
Lagged share of black workers in the department	-0.039
	(0.017)
Manager FE	Y
Store FE	Y
Hiring month and year FE	Y
Outcome mean	0.285
Observations	110,449

NOTE. Clustered standard errors at the manager level are presented in parentheses. See Table 2 for details.

Appendix D Additional groups

Table D1: OLS estimates of the cumulative impact of previous experiences on current hiring, negative and positive experiences

Black or Hispanic fractiong hired	Black managers (1)	Female workers (2)	Hispanic workers (3)	Black managers (4)	Female workers (5)	Hispanic workers (6)
Black fraction quit/fired $\leq$ 3 months	-0.266	-0.057				
	(0.054)	(0.022)				
Hispanic fraction quit/fired $\leq$ 3 months			-0.025			
			(0.019)			
White fraction quit/fired $\leq$ 3 months	0.084	0.031	0.029			
	(0.072)	(0.027)	(0.023)			
Black fraction tenure $\geq$ 12 months				0.203	0.072	
				(0.081)	(0.031)	
Hispanic fraction tenure $\geq$ 12 months						0.003
						(0.024)
White fraction tenure $\geq$ 12 months				0.048	-0.047	-0.016
				(0.074)	(0.038)	(0.027)
Manager FE	Y	Y	Y	Y	Y	Y
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y	Y	Y
Outcome mean	0.556	0.403	0.293	0.551	0.402	0.290
Standard deviation black	0.237	0.291	0.294	0.210	0.2221	0.253
Standard deviation white	0.240	0.232	0.205	0.225	0.205	0.209
Observations	3,396	19,546	27,349	2,825	16,198	22,482

NOTE. Clustered standard errors at the manager level are presented in parentheses. See Table 2 for details.