Swedish Institute for Social Research (SOFI)

Stockholm University

WORKING PAPER 10/2023

LEARNING TO DISCRIMINATE ON THE JOB

by

Alan Benson and Louis-Pierre Lepage

SAMMANFATTNING AV SOFI WORKING PAPER 10/2023

När minoritetsgrupper ansöker om jobb eller interagerar med rättssystemet kan utfallen komma att bero på snedvridning (”bias”) hos inflytelserika beslutsfattare. Det görs avsevärda ansträngningar för att minska risken för bias bland samhälleliga institutioner som fattar viktiga beslut. Dessa ansträngningar innefattar bl.a. mångfald inom beslutsgrupper, algoritm-baserade rekommendationer, eller kurser för att motverka bias. Men dessa åtgärder kan bara lyckas om beslutsfattares uppfattningar är formbara och inte om uppfattningarna är svåra att påverka.

Vi finner att beslutsfattares snedvridna uppfattningar formas av personliga erfarenheter under den specifika anställningen, snarare än att de är förutbestämda och fixerade över tid.

Vi undersöker en serie av anställningsbeslut fattade av 27 470 chefer vid fler än 4 000 butiker inom detaljhandeln i USA. I vår kontext har chefer inom detaljvaruhandeln en avsevärd autonomi när det gäller anställningsbeslut för sina respektive avdelningar. Vi finner att, även inom en given butik, tenderar olika chefer att anställa olika typer av arbetare med avseende på deras etniska bakgrund. Vi undersöker om dessa skillnader grundar sig i chefernas egna erfarenheter när det gäller att anställa arbetare med viss etnisk bakgrund. Som exempel karaktäriserar vi en händelse där en chef ger fastanställning till en arbetare, och denna sedermera slutar eller får sparken inom två månader som en negativ händelse. Vi undersöker därefter hur detta påverkar chefens benägenhet att anställa nya arbetare med samma etniska bakgrund.

Vi finner att negativa anställningsupplevelser som involverar afro-amerikanska arbetare leder till en minskning i benägenheten hos den enskilda chefen att anställa afro-amerikaner i framtiden. Detta mönster tycks drivas av några enskilda faktorer: till att börja med tycks chefernas uppfattning om afro-amerikanska arbetares prestationer, både positiva och negativa, ändras oftare än ge gör för vita amerikaner. Dessutom ändrar chefernas erfarenheter deras sannolikhet att anställa afro-amerikaner och vita amerikaner, men återgången till det normala är långsammare för afro-amerikaner. Det här beror på att negativa upplevelser med afro-amerikanska arbetare skapar en negativ perception hos cheferna som gör det mindre sannolikt att arbeta tillsammans framöver, vilket innebär att det tar längre tid att korrigera snedvridningen. Dessa två faktorer samvarierar så att bias gentemot afro-amerikaner är större och mer persistenta än de mot vita arbetare.

Resultaten bidrar till vår förståelse för uppkomsten av bias, vilka verkar uppstå genom individuella upplevelser och utvecklas över tid. Minoriteter är missgynnade av det faktum att chefer tillskriver dem stereotypa egenskaper baserat på de få individer cheferna tidigare har interagerat med, vilket leder till en negativ perception av gruppen som helhet. Vår studie visar också att spridning av information och ett mer utspritt beslutsfattande kan minska risken för bias som påverkar anställningsbeslut.

SOFI Working Paper Nr. 10/2023, ”Learning to Discriminate on the Job”, är författat av Louis-Pierre Lepage, SOFI, och Alan Benson (University of Minnesota). Kontakt: louis-pierre.lepage@sofi.su.se.

Learning to Discriminate on the Job

Alan Benson
University of Minnesota

Louis-Pierre Lepage
Stockholm and Queen's University

April 2023

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 negative and positive experiences with black hires seed the race of future hires, consistent with managers updating their beliefs about the productivity of worker groups. Second, experiences with black workers have a larger impact on future hiring than those with white workers, consistent with greater updating about their productivity. Third, early negative experiences with black workers yield particularly large and persistent declines in a manager's subsequent black hiring, consistent with negative perceptions being slow to correct. Our results suggest that managers' perceptions of worker groups evolve from their individual experiences in a way that systematically disadvantages minorities in the hiring process.

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, Pierre Deschamps, Steve Lehrer, Matthew Lindquist, Erik Lindqvist, Conrad Miller, David Miller, José Montalbán Castilla, Rosina Rodriguez Olivera, Dan-Olof Rooth, Jeff Smith, Aaron Sojourner, and Evan Starr for helpful discussions and feedback as well as seminar participants at the 2022 NBER Summer Institute, ASSA, EEA, SOLE, CIREQ, EALE, CEA, RES, AFSE, ESPE, SEC, the University of Michigan, the Stockholm School of Economics, SOFI at Stockholm University, the Discrimination and Diversity Workshop, 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 unemployment gap has remained approximately constant for decades (Lang and Lehmann, 2012). Recent studies that have sought to understand the root causes of discriminatory behaviors and outcomes have largely yielded a common refrain: that evaluators, judges, physicians, and other influential decision-makers across a variety of domains hold biased beliefs about differences between groups (Reuben et al., 2014; Arnold et al., 2018; Bohren et al., 2019; Bordalo et al., 2019; Sarsons, 2019; Bohren et al., 2021). Moreover, the cumulative effects of everyday decisions informed by biased beliefs could also provide a foundation for understanding systemic discrimination and its consequences for social and economic inequality (Darity and Mason, 1998; Bohren et al., 2022). In this paper, we ask: where do biased beliefs come from and how do they evolve over time?

Using administrative data from one of the largest employers in the US, we investigate whether discriminatory beliefs emerge specifically from a decision-marker's personal experience with their environment, rather than information that is not lived. We posit that a manager's experience hiring workers of different races shapes their beliefs about each race's productivity and therefore influences their subsequent hiring. Moreover, because managers may avoid hiring workers of a given race following bad experiences, negative biases could be particularly slow to self-correct, yielding average beliefs that are negatively-biased against minorities (Lepage, 2022). This makes the emergence of discriminatory beliefs consistent with "experience effects" documented in employer surveys (Pager and Karafin, 2009) and outside of the employment setting (Malmendier, 2021a,b). Unlike most theories of discrimination, we study how biases about groups of workers evolve in predictable ways within a manager and over time. An unusual feature of the discrimination we study is that differences in biased beliefs across managers may not be due to some inherent trait of the manager, but rather the "luck of the draw" among early experiences with worker groups.

Although identifying the causes and consequences of individual biases would help organizations and policymakers design effective remedies, data limitations have stymied efforts to study how individuals' discriminatory beliefs form and evolve (Charles and Guryan, 2011; Guryan and Charles, 2013). For instance, Census and audit study data typically lack information on hiring managers, inhibiting the ability to study the emergence and evolution of individual-level discrimination over time.

We seek to overcome this challenge using data from the US operations of a large retailer which allows us to examine how managers' past experience hiring workers of different races affects the race of their subsequent hires. The data include over 1 million workers in permanent positions working for over 27,000 store managers across over 4,000 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 proceeds by examining whether this manager-level variation in hiring reflects belief formation "on the job" about workers from different races. Because we do not directly observe manager beliefs, we infer belief-updating using within-manager variation in the race of hires, conditional on whether the manager had previous negative or positive experiences hiring different racial groups. To operationalize negative and positive experiences, we use idiosyncratic variation in 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. Informed by our institutional setting, our main analysis classifies positive experiences as new hires who achieve at least 12 months of tenure in their position and negative experiences as new hires who are fired or quit within 3 months. To validate tenure as a measure of good or bad experiences, we show that longer realized tenure is correlated with greater objective sales performance (where this metric is available) and a greater likelihood of leaving for voluntary reasons rather than being dismissed for poor performance, but uncorrelated with store, department, or market-level factors.[1]

[1]Moreover, we document similar effects across alternative experience measures, including a continuous measure of expected tenure based on a hazard-rates approach or a classification of

We have three main findings: (1) a manager's past hiring experiences (positive and negative) affect their probability of hiring from that group in the future, (2) this effect is much more pronounced for black workers than for white workers, and (3) negative experiences with black workers have more persistent effects than positive experiences or negative experiences with white workers. We document that managers respond particularly strongly to early negative experiences with black workers, with estimates from bootstrapped samples suggesting that managers' first experiences alone reduce their subsequent black hiring by 7.3%, from 25% to 23.3% of total new hires at the firm, corresponding to 1,000 black workers each year. We document these findings even though the distribution of negative and positive experiences with black and white hires is similar at the firm. Our findings are consistent with the proposition that managers condition their hiring decisions on their personal experiences on the job, and not with the proposition that beliefs or biases are immutable by the time they become managers.

After establishing our main results, we conduct supplementary analyses examining how managers learn to discriminate on the job. First, we find that hiring experiences have the greatest effect on future hiring early in a manager's career, when they should be most uncertain about the performance of worker groups. Second, if previous experiences of a manager shape their group perceptions, then workers who overcome a negative bias—that is, workers hired by a manager who had negative experiences with their group—should be positively selected. Consistent with this, and in contrast with several alternative mechanisms that could generate persistence in group hiring within managers, we find that a higher share of negative (positive) experiences with previous hires predicts a lower probability of a negative (positive) experience with current hires. Third, we examine other sources of information that managers may be able to draw upon and find little evidence that learning occurs based on hires at other departments within a store.

Our final set of analyses situates manager-specific experience effects against other mechanisms that could yield a path-dependent trajectory in the race of hires. One broad class of alternatives attributes the race of future hires to the hiring history of the department or the manager's team rather than the manager themselves. For instance, referrals, productive complementarities, or employee discrimination could explain why the race of a department's past hires predicts future hires. We present results showing

managers into quintiles based on their cumulative race-specific turnover.

that time, manager, and team composition placebos fail to predict the race of future hires, whereas the hiring experience of the "true" manager adds substantially to our ability to predict future hiring. These analyses require alternative explanations to be manager-specific, such as stationary taste-based discrimination or pre-existing manager biases that could affect both hiring decisions and hiring outcomes.[2] However, these alternatives are not consistent with several of our findings and placebo tests, including documenting that the race of hires depends on whether early experiences are positive or negative, or that a higher share of negative (positive) experiences with previous hires predicts a lower probability of a negative (positive) experience with current hires. We caution that our tests do not rule out that these other factors may be at play, but simply that we identify novel variation in hiring that cannot easily be ascribed to these factors. We propose that accounting for on-the-job hiring experience of individual managers adds to our ability to predict hiring discrimination.

To contextualize our results, we propose a statistical discrimination framework in which managers are initially uncertain about the productivity distribution of different groups and update their beliefs from personal experience with individuals. Unlike classical statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977), the model features managers who update their beliefs about different groups "on the job" through hiring and observation. Like the literature on employer learning (Farber and Gibbons, 1996; Altonji and Pierret, 2001), our model features belief updating from observing individual workers, but crucially, managers in our model also use these observations to update beliefs about *groups* rather than only specific individuals. As a result, managers hire more from groups with whom they have positive experiences and less from groups with whom they have negative experiences, even if these positive or negative experiences are idiosyncratic.

Our framework offers a parsimonious explanation for our findings based on greater updating about minority groups and endogenous learning: positive experiences beget faster hiring and learning, whereas negative experiences slow hiring and learning.[3] Therefore, the initial "seeding" of hiring experiences with minorities can have

[2]Similarly, differences in applicant pools or worker sorting across managers based on their experiences with a worker group provide poor alternative explanations for our findings, as discussed in Section 6.

[3]The key claim of our model is that managers update beliefs about groups from individual experiences. Greater updating about minority groups may arise due to weaker priors or behavioral biases (Allport et al., 1954; Bordalo et al., 2016, 2023).

substantial, persistent effects on hiring discrimination, systematically reducing black hiring across managers at the firm. Intuitively, if experience is a product of a decision-maker's own decisions, then they can become insulated from new information that would debias their beliefs.

Our findings offer an explanation for why economic gatekeepers in a wide variety of domains hold persistently biased beliefs about minority groups. For instance, Bohren et al. (2019) finds that evaluators hold biased beliefs about the productivity of women in a online market, which may be overcome at the individual level through additional performance information. Instead, we consider discriminators who update beliefs about *groups* based upon information on individuals. Closer to our paper, Sarsons (2019) finds that physicians are less likely to refer a patient to a female surgeon following a negative experience with a female surgeon, but the same is not true for male surgeons and the opposite is not true for a positive experience. We also find that minority workers are disproportionately disadvantaged in the hiring process when managers have had more negative experiences with their group. However, hiring discrimination in our context best fits within a dynamic learning process in which information acquisition is endogenous, rather than a static bias in belief updating (e.g. attribution bias) as in their context.[4] Another set of studies attributes differences in economic outcomes to implicit biases, also at the level of the manager (e.g. Glover et al., 2017), though it is unclear whether these biases have informational roots or evolve over time. We propose and document that biased beliefs should not be thought of as static; they arise and evolve specifically through personal experience.

We also contribute to a large literature on racial discrimination and inequality. We study one of the largest employers of low-skill workers in the US, jobs for which racial discrimination is often posited to be largest and could differ importantly from gender discrimination or discrimination against high-skill workers (Lang and Lehmann, 2012). Recent evidence from correspondence studies indicates hiring discrimination at several large US employers (Kline et al., 2022). We take a complementary approach by documenting a specific source of racial discrimination at one such large employer.

[4]Some of our analyses document an asymmetry in manager responses following positive versus negative experiences, but our findings are not inconsistent with Bayesian updating by managers over their own experiences and attribution bias is not sufficient to explain our results. The updating process is important because policy implications can differ sharply. It is unclear whether additional information or experiences would mitigate discrimination if it is driven by a static behavioral bias. In contrast, our evidence is consistent with a framework in which additional experiences promote learning and correct negative biases.

Such jobs are also often a first rung in a career ladder, with important implications for future outcomes and the persistence of inequality (Oreopoulos et al., 2012). We present the first evidence that individual manager biases arise from experience and create systematic racial discrimination in hiring, generating novel implications for understanding the roots of racial inequality.

Our findings 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; Li et al., 2020). A literature documents heterogeneous hiring discrimination across managers, but we focus on their own market interactions with groups rather than differences by manager race or gender (Giuliano et al., 2009; Giuliano and Ransom, 2013; Åslund et al., 2014; Hjort, 2014; Glover et al., 2017; Cullen and Perez-Truglia, 2021; Ronchi and Smith, 2021). Our results relate specifically to the literature on "experience effects" (Malmendier and Nagel, 2011; Leung, 2018; Malmendier, 2021a,b). 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 influenced by experience effects that operate similarly in labor markets, providing a new explanation for persistent between-group outcome disparities. More generally, our results provide novel evidence regarding the trade off that firms face between extraction, in this case hiring from a majority group with better known productivity, and exploration, hiring from a minority group with more uncertain productivity. This type of trade-off has long been recognized as fundamental to organizational design (March, 1991; Denrell and March, 2001), relating to the economics literature on bandit problems (Bergemann and Valimaki, 2006).

Our results give cause for both pessimism and optimism for efforts to combat discrimination. On 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 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, such as hiring algorithms or affirmative action policies that may accelerate learning and the correction of biased beliefs (Miller, 2017; Paluck et al., 2019; Li et al., 2020).

The remainder of the paper is organized as follows. Section 1 describes our data and institutional setting. Section 2 presents our empirical approach. Section 3 presents our main results on how previous experiences of managers with groups generate hiring discrimination. Section 4 presents additional results regarding employer learning from experience. Section 5 investigates alternative interpretations of our main results. 6 interprets our findings using a simple theoretical framework. Section 7 concludes.

1 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 (age, gender, race), job, department, and location. We also observe employment termination including dismissals, quits, and layoffs. Each store is led by one store manager and a set of department managers who hire for their respective department (on average 4-5 managers 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 they are the two largest racial categories in our data, which makes it most feasible to estimate managers' evolving hiring behavior.[5] Summary statistics on workers and managers are presented in 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 influence their hiring, we consider a manager-level panel in which one observation corresponds to a month in which

[5]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.

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)	0.144	(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.

a manager hires at least one black or white worker, which we refer to as a hiring event. Our analysis restricts our sample to new managers who began hiring for the first time during our sample period, since we can observe their first hiring experience with workers at the firm. We study managers who were at least willing to hire and manage black workers over our sample period, potentially excluding some managers with strong initial bias, but our focus is on whether and how bias evolves over time. 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 concentrated in periods managers hire workers and arises irregularly across time.[6]

[6]In these specifications, we control for the number of hires in a given hiring event and the time between hiring events to account for the possibility that negative and positive experiences affect

Staffing levels for permanent positions in each store are determined by forecasts made by the firm's headquarters. When a manager is tasked with filling a vacancy, the manager would typically begin by requesting a shortlist of candidates from the location's 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 to 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 a HR representative of their interview performance. Anecdotal evidence from store managers we interviewed indicates that it would be very rare for a manager to have prior familiarity with a new applicant, though there is no formal process for tracking referrals and we do not observe these instances directly. More generally, as we discuss below, there is limited room for individual managers to influence the applicant pool they receive for a position, and applicant pool endogeneity, sorting between workers and managers, or referral hiring provide poor alternative explanations for our findings.

Turnover at the firm is high, in line with the retail sector more generally which has 50% greater turnover than the US average.[7] High turnover provides valuable variation in hiring of workers within managers even over 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 around 10 weeks of worker salary (Boushey and Glynn, 2012). Hiring and retaining high quality workers into permanent positions is accordingly one of the most important aspects of department managers' jobs.

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

the time in between hiring events, but these controls have little impact on the results. Focusing on a manager's own experiences also ignores potential spillovers from experiences across managers. Ultimately, results shown below indicate that managers' updating behavior is qualitatively similar whether considering a time or a manager-event panel and that managers do not seem to learn from the experiences of other managers at their store.

[7] https://www.bls.gov/news.release/jolts.t18.htm.

Figure 1: Cumulative turnover by tenure

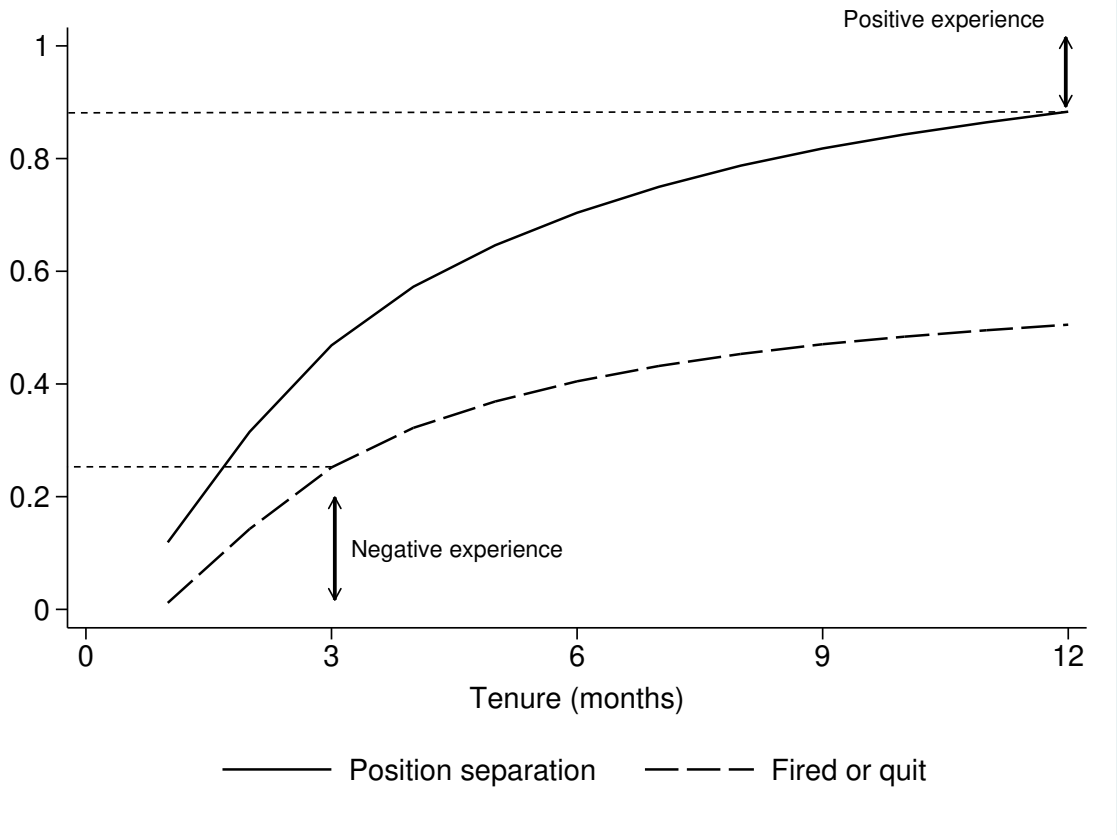

NOTE. "Position separation" refers to the worker no longer working in the position for which they were hired, including dismissals and quits but also department or store transfers, layoffs, promotions, and retirement/disability/death. The sample is restricted to workers hired at least one year before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

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, turnover declines substantially and remains below 2% per month.

There is large variation in black hiring across managers at the firm. 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 hires is 56% and the median is 59%. Appendix B presents evidence that variation in the share of black hires persists after controlling for the manager race as well as store, department, and job effects. A substantial share of residual variation is explained by manager fixed effects,

implying that individual manager factors play a substantial role in determining the race of hires. Next, we study whether these individual effects are in fact dynamically driven by experiences on the job.

2 Belief updating from experience

Our main empirical analysis examines whether manager-level variation in hiring can be explained by managers' updating from their own prior experiences. Our first specification investigates how cumulative experiences with black and white workers in previous hiring events affect the share of black hiring (restricted to black and white hires) in the current event. We estimate the model

$$FB_{emlt} = \beta_1 \overline{EXP}_{B,e-1} + \beta_2 \overline{EXP}_{W,e-1} + X_{emlt}\zeta + \theta_t + \lambda_l + \gamma_m + \varepsilon_{gemlt} \qquad (1)$$

where the dependent variable is the fraction of black workers hired in hiring event e by manager m in location l at time t. $\overline{EXP}_{B,e-1}$ and $\overline{EXP}_{W,e-1}$ respectively indicate the share of black and white hires up to event $e-1$ for whom the hiring experience was negative, as defined below, and their coefficients reflect how negative hiring experiences affect the race of hires in the current event. Other specifications consider indicators for positive, rather than negative, experiences. 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. Time fixed effects account for potential differences in the applicant pool and worker performance at the firm over time. 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, although results are similar when clustering at the store level.

We use the coefficients on $\overline{EXP}_{B,e-1}$ and $\overline{EXP}_{W,e-1}$ to test our key predictions. Intuitively, near-zero estimates suggest that heterogeneity in the race of managers' hires can be explained by factors relating to the hiring context and (potentially

unobserved time-invariant) manager characteristics. In contrast, if past experience predicts 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, negative and positive experiences with a group would appear to affect managers' beliefs and therefore their propensity to hire from a group.

While this specification provides an intuitive way to investigate how an employer's hiring history affects the race of current hires, interpretation can be complicated because the decision to hire from a group as well as the quality of hiring experiences after the first can be endogenous to previous experiences, as would be the case if managers adjust their hiring thresholds based on past experience. For instance, if a manager has a bad experience with a black worker then sets a higher bar for hiring the next black worker, but still has another negative experience, then the manager may update their beliefs more after their second hire given that they performed poorly despite overcoming a greater hiring bar.

Therefore, our analysis proceeds by restricting the sample to a manager's first hiring experience only, or their experience with incumbent workers

$$FB_{emlt} = \beta_1 EXP_{B,1} + \beta_2 EXP_{W,1} + X_{emlt}\zeta + \theta_t + \lambda_l + \varepsilon_{gemlt} \tag{2}$$

where $EXP_{B,1}$ and $EXP_{W,1}$ respectively indicate the share of black and white hires in the manager's first event with each race for whom the hiring experience was negative or positive. These specifications test how a manager's first experience affects their subsequent hiring, potentially setting them on persistently different belief-updating paths. Moreover, we can test whether a manager's first hiring outcome with black workers is exogenous to initial department conditions after accounting for store-level factors, providing us with a source of exposure to worker groups which is plausibly uncorrelated to factors outside of the manager themselves.

2.1 Measuring negative and positive experience

Because our analysis considers hiring decisions as a function of idiosyncratically negative or positive experiences, we must distinguish hiring events as either negative or positive versus a manager's expectation. Since we observe some performance measures, but not their discrepancy with a manager's expectation, we use different performance measures relative to other new hires at the firm, which should inform

expectations.

Our main specification identifies particularly negative and positive experiences of a manager. We measure the degree to which a manager's hiring experience with a given race is negative by calculating the share of new hires of that race who were fired or quit in the first 3 months of employment. As shown in Figure 1, the first three months represent a key period after which the rate of dismissal and quit decreases substantively.[8] Workers hired into permanent positions who leave or who are terminated within the first three months account for around a quarter of hires. Such turnover is very costly: workers must be hired, trained, and provided time to develop tacit skills and a familiarity with the store's protocols and products. Workers who depart after short tenures also impose an opportunity cost: they filled a spot that could have otherwise been filled by a successful hire.

We measure the degree to which a manager's hiring experience with a given race is positive by the share of hires of that race who achieved at least 12 months of tenure in the position for which they were hired. As shown in Figure 1, after 12 months, the probability 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 be transferred or laid off. Approximately 15% of new hires achieve tenure of at least one year in their position.

We do not posit that the positive experience is only revealed to managers after one year. Rather, we use this ex-post measure as a way for us to characterize hires which were successful in the absence of direct shorter-term performance information. We expect that managers in the first few months of employment already observe a positive signal of quality for a worker who will eventually achieve over one year of tenure. Consistent with this conjecture, some jobs in our sample feature performance metrics based on sales figures and Online Appendix A shows that, within this subset of sales-commissioned workers, workers who will eventually achieve at least one year of tenure clearly outperform workers who quit or are fired over the first three months.

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 (27% vs

[8]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 (Table C3).

25%) and lower 12-month retention rate (12% vs 14%). Most variation in tenure is not explained by race, but by idiosyncratic differences across individuals within race. For instance, there's a 48% probability that a given black hire meets or exceeds the average tenure of a white hire. Black and white workers have similar median (4 versus 4.5 months) and average tenures (5.6 versus 6.1 months). Moreover, conditional on being fired or quitting within 3 months, both black and white workers achieve an average tenure of 2.4 months, so there is little difference in tenure across race conditional on an experience being classified as negative. While our proposition that managers adjust their group perceptions and hiring discrimination with experience does not depend on whether the performance of black and white workers is the same, it is straightforward to think of managers drawing from two similar performance distributions but reacting differently to the draws they observe from the black distribution. Moreover, 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 these managers endogenously avoid hiring black workers in the future.

By using tenure to classify negative and positive experiences, we are not asserting that objective worker performance is the only factor affecting hiring and retention. Differences in tenure across race, though modest, 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 (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 a manager's personal, lived 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.

Even though managers should directly value tenure in a position, we would ideally also have a direct measure of productivity that the manager observes. 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 (Bycio et al., 1990; Williams and Livingstone, 1994; Huang et al., 2006; Zimmerman

and Darnold, 2009). Further, Online Appendix A presents additional evidence that the typical performance-tenure relationship holds for sales-commissioned positions. Another caveat when using tenure is that positive and negative experiences could be partially determined by retention efforts of managers which correlate with hiring efforts, for example through business cycle shocks. In section 5, we present evidence that differential retention efforts across worker race are not driving our results.

3 Main results

First, we test whether a manager's experiences hiring workers influence their subsequent hiring of workers from the same racial group. Table 2 presents estimates of the relationship between negative and positive previous experiences and the race of 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 effects in percentages are approximately 50% larger for black than white hiring given that they constitute a minority of workers, indicating that hiring experiences play a disproportionately large role in black hiring. Unless specified otherwise, all of our tests refer to a statistical significance level of 5%.

The first three columns of Table 2 present estimates from equation (2) indicating that managers significantly decrease their hiring of black workers by an estimated 6% in column 3 for a one standard deviation increase in the fraction of previous black hires that were dismissed or quit within 3 months. Estimated impacts for experiences with white workers indicate a substantially smaller but still statistically significantly increase of approximately 3% in black hiring in column 3, when accounting for the higher standard deviation of experience measures with black workers.

Columns 4-6 present estimates of the impact of positive experiences. Managers significantly increase their hiring of black workers by an estimated 4% in column 6 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 not statistically significant at conventional testing levels.

The key takeaway from Table 2 is that a manager's propensity to hire black workers in the future depends on whether their experiences hiring black workers were

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

Black fraction hired	(1)	(2)	(3)	(4)	(5)	(6)
Black fraction quit/fired $\leq$ 3 months	-0.053	-0.073	-0.072			
	(0.009)	(0.016)	(0.017)			
White fraction quit/fired $\leq$ 3 months			0.044			
			(0.022)			
Black fraction tenure $\geq$ 12 months				0.029	0.057	0.058
				(0.014)	(0.022)	(0.024)
White fraction tenure $\geq$ 12 months						-0.001
						(0.027)
Manager FE		Y	Y		Y	Y
Worker and event controls		Y	Y		Y	Y
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Outcome mean	0.367	0.367	0.348	0.366	0.366	0.347
Standard deviation black	0.295	0.294	0.294	0.229	0.229	0.229
Standard deviation white	0.189	0.188	0.188	0.205	0.205	0.205
P-Value: B = -1 * W			0.060			0.090
Observations	34,496	33,971	31,911	28,879	28,456	26,655

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. When indicated, 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, as well as month and year, store, and manager fixed effects. We exclude workers hired in the last 3 months (1 year) of our sample period for negative (positive) experiences since we cannot compute experience measures for these workers.

positive or negative in the past. In the interest of simplicity, our preferred specification adopts a number of design choices: updating occurs at each hiring event, we discretize positive and negative experiences rather than a continuous measure or quartiles, and we treat short-term quits and terminations as equivalently negative. Further analysis finds that our results are not sensitive to these specification decisions, though we do find larger effects for terminations than quits (Tables C1, C2, C3).[9]

[9]Table C1 presents results organizing the data into a time panel and considering a continuous performance measure comparing tenure achieved by a manager's hires to expected tenure at the firm. Specifically, using a hazard rate approach, we compute deviations in turnover rates by race and months of 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. 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. Results are qualitatively and quantitatively similar comparing to race-specific turnover rates or average turnover rates across racial groups. Table C2 defines positive (negative) experiences as the previous hires of a manager being in the top (bottom)

As discussed in more detail in Section 5, updating from negative experiences across groups is similar in periods of high versus low labor market tightness as measured using unemployment rates, which may influence the relative quality of outside options across race, suggesting that our results are not driven by differential turnover reasons across race (Table C8). Lastly, results are qualitatively similar when restricting to female workers or black managers, suggesting that both black and white managers respond similarly to their previous experiences with black workers (Table D1).

Next, we focus on first experiences to avoid conditioning on an endogenous sequence of hiring experiences, showing instead that first experiences can set employers on persistently different paths of belief-updating and hiring. Initial experiences may be particularly salient because managers presumably have weaker priors, but a first experience alone can be quite misleading regarding a group's expected productivity. Indeed, while the distribution of realized tenure among first hires is very similar across race, realized tenure of first hires varies widely within race (Figure C1). Still, managers seemingly update quite strongly from a first negative experience with black workers, as we show next.

We begin by estimating equation (3) and plotting the results in Figure 2, showing the persistent effect of first hiring experiences on the race of subsequent hires. As shown in the top left panel, when a manager's first black hire(s) quit or are fired within three months, the manager is persistently less likely to hire black workers over the next 6 hiring events, corresponding to around 15 months on average.[10] The decrease is strongest immediately following the negative experience, corresponding to a reduction of around 14% versus 5% for events 4-6. The magnitude and persistence of this effect is specific to negative experiences with black workers, rather than white workers or positive experiences. In particular, while positive experiences temporarily increase black hiring as we will show below, the impact of a first positive experience appears to have dissipated within the following three hiring events, corresponding to approximately 8 months. Since 0.7 black workers per event are hired on average, these results imply that manager perceptions following a positive experience have reverted

quartile of deviations from expected tenure at the firm also yields similar results. Table C3 separating firings from quits indicates that firings lead to larger negative impacts, but that managers respond negatively to both.

[10]Restricting the sample to managers with at least 10 hiring events yields similar results over the hiring events that immediately follow, but the impact seems less persistent for later hiring events when excluding managers who hire fewer times in total.

Figure 2: Impact of a manager's first hiring experience with black and white workers on their subsequent black hiring share

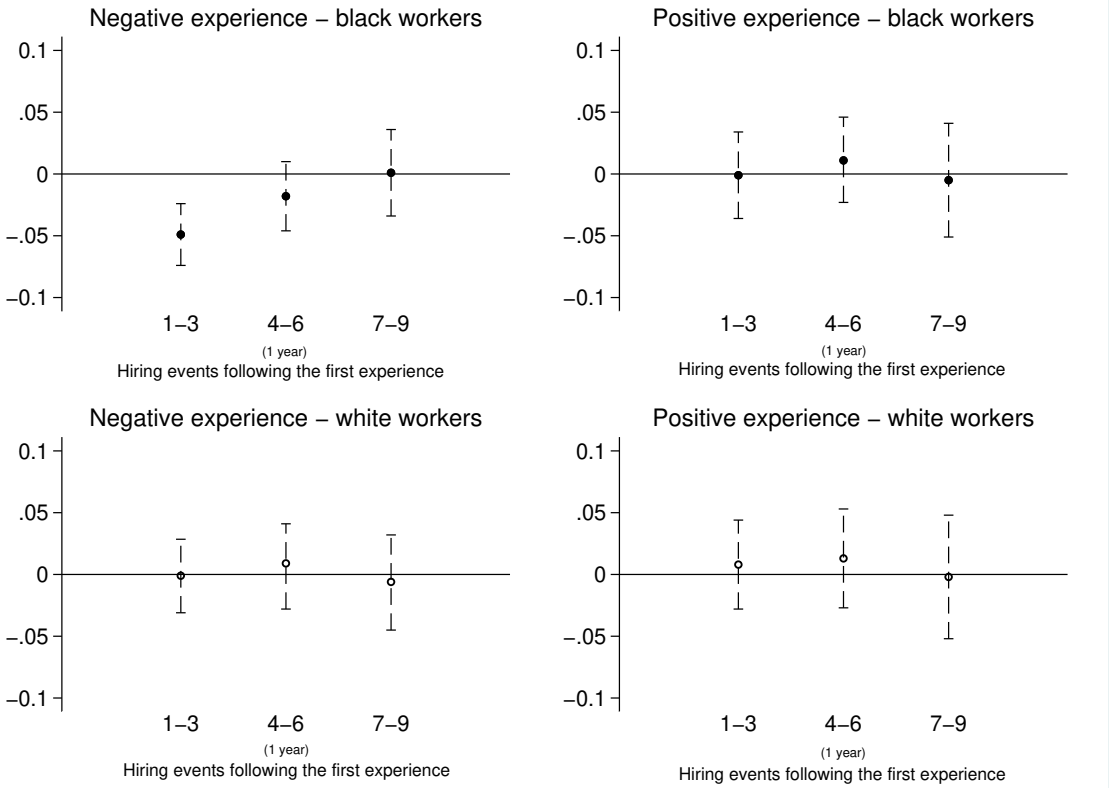

NOTE. 95% confidence intervals from clustered standard errors at the manager level are presented using dashed lines. A negative experience corresponds to the fraction of a manager's hires from a racial group in their first hiring event that was fired or quit in the first 3 months of employment. A positive experience corresponds to the fraction of a manager's hires from a racial group in their latest hiring event achieving 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. We exclude workers hired in the last 3 months (1 year) of our sample period for negative (positive) experiences since we cannot compute experience measures for these workers.

after having observed slightly more than two new black hires on average.

Figure 3 shows that a first negative experience with black workers sets managers on a persistently different hiring path. The fraction of managers who hire at least one black worker in the hiring event following their first experience is about 15 percentage points lower at around 33% for managers whose first experience was negative. The hiring gap subsides as managers hire more workers, but there remains a 5 percentage points difference in the fraction of managers who have hired at least one black worker

Figure 3: Share of managers who have hired at least one black worker following their first experience with black workers

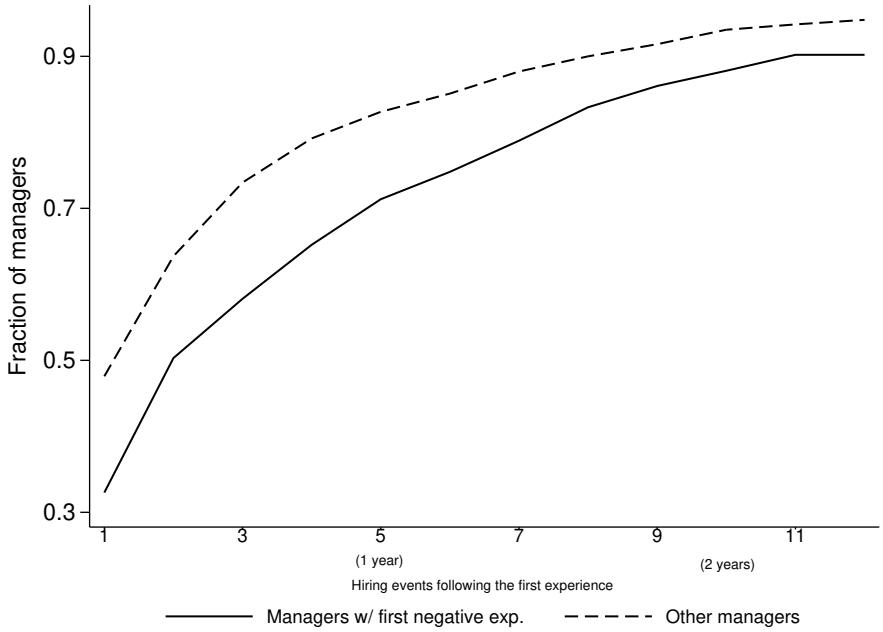

NOTE. See Figure 2 for additional details.

even after 12 hiring events following their first experience, corresponding to a period of around 2.5 years.

Similar results are shown in Table 3. Estimates from the first three columns indicate a statistically significant decrease of 2% in black hiring in the current event for a one standard deviation increase in the fraction of the first black hire(s) that were fired or quit within 3 months. These results suggest that early negative experiences with black workers impact hiring over our entire sample period. Columns 4-6 show smaller statistically non-significant impacts for early positive experiences, suggesting that the impact of early positive experiences does not persist, consistent with Figure 2. Impacts of a first negative experience with white workers are smaller, statistically non-significant, and we can reject that the coefficient is of the same magnitude but opposite sign to that of a first negative experience with black workers.

One potential concern with our analyses would be that our specification's store and time fixed effects do not capture department-specific factors affecting the performance and hiring of black workers through mechanisms unrelated to the manager. To

Table 3: OLS estimates of the impact of a manager's first experience with black and white workers on current black hiring, negative and positive experiences

Black fraction hired	(1)	(2)	(3)	(4)	(5)	(6)
Black fraction quit/fired $\leq$ 3 months	-0.024	-0.021	-0.019			
	(0.007)	(0.007)	(0.007)			
White fraction quit/fired $\leq$ 3 months			-0.009			
			(0.008)			
Black fraction tenure $\geq$ 12 months				-0.010	-0.007	-0.010
				(0.010)	(0.010)	(0.011)
White fraction tenure $\geq$ 12 months						-0.003
						(0.011)
Worker and event controls		Y	Y		Y	Y
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Outcome mean	0.410	0.387	0.367	0.411	0.387	0.366
Standard deviation black	0.405	0.405	0.405	0.327	0.327	0.327
Standard deviation white	0.348	0.348	0.348	0.342	0.342	0.342
P-Value: B = -1 * W			0.006			0.375
Observations	39,143	36,816	35,613	32,969	30,908	29,869

NOTE. Robust standard errors are presented in parentheses. See Table 2 for additional details.

examine this possibility, we run a set of placebo tests to see if our results are reproduced at counterfactual first hires. First, we examine whether performance of black workers *before* the manager begins indicates the race of the manager's first hires; we find no significant evidence this is the case. Second, we examine whether the fraction of black hires before a manager begins predicts the performance of a manager's first black hires; again, we find no evidence this is the case (Table C6). Another potential concern is that early negative experiences result from pre-existing bias which is also associated with lower black hiring, for example due to taste-based discrimination, rather than experiences themselves shaping group perceptions. To examine this possibility, we run a placebo test to see if lower hiring of black workers by a manager in their first event, which under this alternative explanation would capture a pre-existing manager bias, predicts a higher likelihood of having negative experiences in subsequent hiring events. We find little evidence that this is the case (Table C6); rather, it is truly early experiences which shape subsequent hiring. Taken together, our placebo tests imply that potentially viable alternative explanations would need not only to coincide with the timing of individual managers hiring their first workers, but also evolve specifically based on the nature of their hiring experiences.

In principle, additional experiences should solidify managers' beliefs. Therefore, we relax the first-hires restriction by examining the first three hiring events (Table C5). More negative experiences over the first three hiring events appear to accumulate into larger negative impacts. More positive experiences also appear to have larger positive impacts, but the coefficients are not statistically significant.

In Tables 2 and 3 as well as across different experience measures considered in our robustness checks, estimated coefficients on experiences with black hires are larger than those with white hires for both negative and positive experiences. Estimated coefficients on black hires are generally over 30-40% larger 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 at the 10% level, as shown by the p-values presented in Table 2. Overall, evidence of updating is weaker and somewhat inconsistent regarding previous experiences with white workers, though coefficients are of the hypothesized sign.

Our results corroborate that managers disproportionately and persistently reduce black hiring following negative experiences, consistent with them developing persistent negative group perceptions. Taken together, they imply that hiring responses of managers following their experiences with workers systematically decrease black hiring.

To put the experience effects we document into context, we consider two approaches. First, we begin from Column 3 of Table 2 and compare the adjusted R-squared of the full specification compared to that of specifications excluding either our two cumulative experience measures or manager fixed effects. The adjusted R-squared of the full specification is 0.414, while the one excluding manager fixed effects is 0.413 and the one excluding experience measures is 0.389. This exercise suggests that our experience measures explain roughly 6% of the variation in black hiring, a much larger share than manager fixed effects do. Second, we estimate the impact of early negative experiences on black hiring compared to a counterfactual scenario in which managers always hire black workers at the rate they do in their first event (25%), before their first hiring experience. Specifically, we compute the sum of the four coefficients in columns 3 and 6 of Table 3 weighted by the frequency of negative and positive experiences with each worker race across 9,999 bootstrap samples. On average, given that managers respond particularly strongly to negative experiences with black workers, we estimate that a manager's first experience reduces

their subsequent black hiring by 7.3% (standard deviation of 0.057, less than 10% of estimates are above 0). Differently put, without these experience effects, we estimate that black workers would compose 25% of the firm's total new hires rather than the observed 23.3%, corresponding to an additional 1,000 new black hires each year. While this exercise hinges on several approximations and assumptions, it suggests that experience effects are quantitatively important, particularly given that our analysis of first experiences would presumably underestimate cumulative impacts.

4 Examining how managers update beliefs

4.1 Updating over a manager's career

Thus far, we have presented evidence that managers' hiring decisions are influenced by their previous experiences, in particular those with black workers. If these impacts reflect evolving group perceptions as managers update their beliefs, then the extent to which they respond to each additional experience should diminish.

We examine this possibility in Table 4, which 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 2, 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 together 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 in the early, middle, and late segments of a manager's hiring history all affect black hiring. They also highlight that the impact decreases with hiring experience: the relationship between experiences and hiring is strongest early in a manager's hiring career and weakest in the last tercile.

Next, we investigate how a manager's most recent experience affects their hiring. Recent experiences may be salient to managers even as they acquire hiring experience, for example if the hiring context changes over time or due to recency bias (Agarwal

Table 4: 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 (0.037)	-0.082 (0.036)	-0.068 (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.

et al., 2008; Gallagher, 2014; Erev and Haruvy, 2016). Figure 4 shows that managers whose latest black hire(s) quit or were fired within 3 months of being hired decrease the share of black workers they hire in their following hiring event by around 3.5 percentage points or 10%. Similar results in regression form are shown in Table C4. Compared to the impact of a first negative experience, the decrease subsides more rapidly over subsequent hiring events and has mostly dissipated after three events, corresponding to around 7.5 months on average. In contrast, managers whose latest white hire(s) quit or were fired within 3 months of being hired increased their share of black hires in their following hiring event, though the effect is almost 50% smaller and seems to subside more quickly. The figure also shows that positive experiences affect subsequent hiring. The impact is again larger (approximately 30% larger) for experiences with black workers, although it appears smaller than for negative experiences and has subsided by the second event following the positive experience.

Overall, our results suggest that managers weight their most recent experience, but within a broader learning process through which the weight they put on their experiences decreases as they hire more workers.

Figure 4: Impact of a manager's latest hiring experience with black and white workers on their subsequent hiring

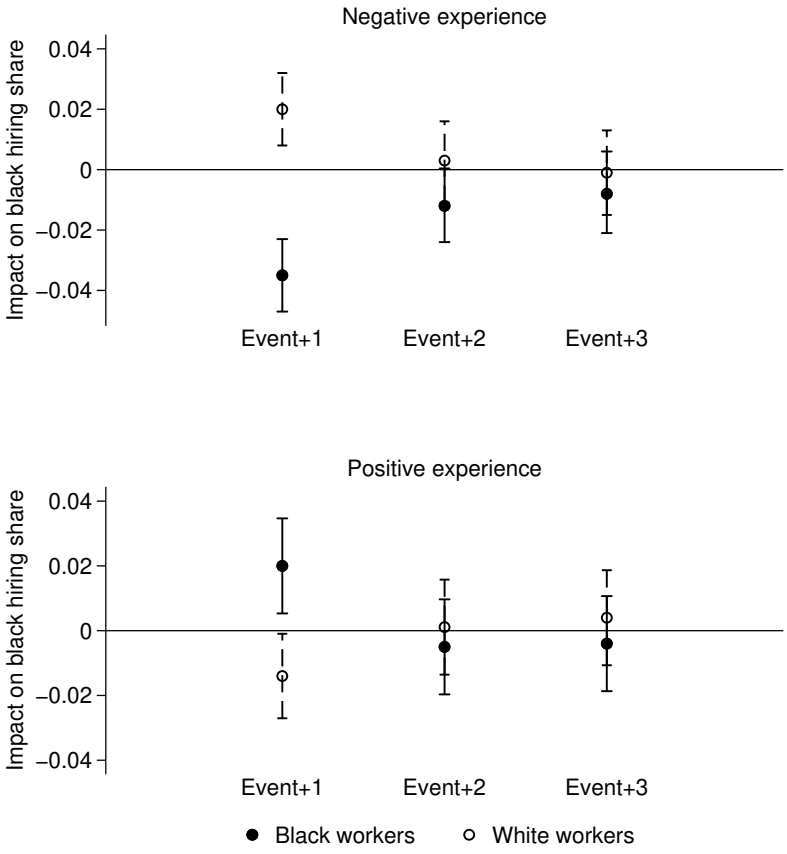

NOTE. 95% confidence intervals from clustered standard errors at the manager level are presented using dashed lines. A negative experience corresponds to a manager's hires from a racial group in the last hiring event being fired or quitting in the first 3 months of employment. A positive experience corresponds to a manager's hires from a racial group in the last hiring event achieving 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, number of previous hiring events, time since the last hiring event, yearly unemployment and college attainment rates in the state, month and year, manager, and store fixed effects. We exclude workers hired in the last 3 months (1 year) of our sample period for negative (positive) experiences since we cannot compute experience measures for these workers.

4.2 Positive selection among workers who overcome bias

If managers develop negative perceptions about the performance of black workers following negative experiences, then they may increase their hiring threshold for subsequent black candidates. In that case, following a manager's negative hiring experiences with black workers, black workers who are hired anyway should be positively-selected and therefore less likely to yield a negative experience. Conversely, following a positive experience, the manager may lower the hiring bar for black workers, decreasing the probability of future positive experiences.

Consistent with this intuition, Table 5 shows a 9% decrease in the probability of a negative experience for a one standard deviation increase in the cumulative fraction of previous negative experiences (4% non-statistically significant decrease for positive experiences). The asymmetry between negative and positive experiences suggests that managers may be prone to reacting more strongly to negative experiences by increasing the bar for hiring black workers, but this pattern could also reflect differences in our measures for good versus bad 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. This pattern is inconsistent with the idea that our main results are driven by mechanisms which simply reflect persistence in a manager's good or bad experiences with black workers that is correlated with the manager's propensity to hire these workers, like hiring through referral, taste-based discrimination, or worker performance that is endogenous to manager bias.

4.3 Geographic heterogeneity

The extent and persistence of decreases in black hiring following negative experiences of managers likely varies based on the relative size of the black workforce in the area. In areas where black workers are a smaller minority group, managers may have less experience interacting with them and weaker priors about their performance. Moreover, the smaller the minority share, the slower the arrival rate of new minority hires, even in the absence of bias, and therefore the slower negative perceptions may dissipate. Lastly, if managers purposefully aim to reduce black hiring following negative experiences, it may be easier to do so in areas where black workers make up

Table 5: OLS estimates of the impact of previous experiences with black workers on the probability of having a negative or positive experience with black workers in the current event

	Fraction quit/fired $\leq$ 3 months (1)	Fraction tenure $\geq$ 12 months (2)
Cumulative fraction quit/fired $\leq$ 3 months	-0.090	
	(0.016)	
Cumulative fraction tenure $\geq$ 12 months		-0.031
		(0.020)
Store FE	Y	Y
Hiring month and year FE	Y	Y
Worker and event controls	Y	Y
Outcome mean	0.221	0.110
Standard deviation	0.225	0.147
Observations	17,000	14,207

NOTE. See Table 2 for additional details.

a small share of applicants.

We test the proposition that negative experiences with black workers have stronger and more persistent effects in areas with a smaller black population. To do so, we match stores to data on racial composition by ZIP Code from the American Community Survey and estimate equation (3) separately for stores located in areas with above and below-median black-to-white population ratios. Consistent with our proposition, negative experiences appear to lead to much larger, more persistent declines in black hiring in areas with smaller black populations. The likelihood of hiring a black worker over the three events following a first negative experience declines by 46% (s.e. 17%) in areas with low black populations versus 13% (s.e. 8%) in those with high black populations. The decrease in low-black population areas remains large at 42% (s.e. 19%) for the next three events and 23% (s.e. 25%) for events 7-9, while it has mostly dissipated by events 4-6 in other areas (Figure C2). Although individual coefficients are not statistically significantly different across areas, the joint test that coefficients on black hiring following a first negative experience are equal in both areas across events 1-9 rejects the null at the 10%

threshold (p-value=0.097).

4.4 Other person-specific sources of belief updating

In addition to their own experience hiring for their department, managers may update their beliefs from different sources: their colleagues, experiences at previous stores, experiences with workers that were already in the department when they joined, and negative experiences less likely to result from a bad match between the manager and the worker (Table 6).

First, managers may update their beliefs from their peers within the store, which could mitigate the extent to which they rely on their own personal experiences. To examine this possibility, we add measures of cumulative negative experiences for other managers within a store to equation (2). The estimated impacts of a manager's own negative experiences with black workers remain largely unchanged, while estimated coefficients on experiences of other managers within the store are smaller and statistically non-significant. Even in a setting where same-store managers' experiences are fairly easy to observe, they appear to have little impact on a manager's hiring after accounting for their own experiences.

Second, manager beliefs may carry over when they move to a new store. To examine this possibility, we estimate equation (2) restricting the analysis to 977 hiring events of managers joining a new store for whom we observe hiring outcomes at both the origin and destination stores. Only 29% of managers from the corresponding analysis presented in Table 2 meet these criteria, placing considerable demands on the data. We find similar point estimates as our main analysis, though errors are outside conventional testing thresholds, providing suggestive evidence that managers make hiring decisions based upon good and bad experiences at their previous stores.

Third, we leverage a source of manager exposure to black workers that is more plausibly exogenous to their hiring decisions. We find that the impact of a manager having a negative experience with "endowed workers" who were already in the department when the manager joined (higher share of incumbent black workers who quit or are fired by the time the manager hires workers for the first time at the department), and as such were not hired by the manager, is similar to the impact of a manager having a negative experience with their first black hire(s).

Our fourth analysis examines whether a negative experience less likely to have

Table 6: OLS estimates of the impact of previous experiences with black workers on current black hiring, other sources of updating

Black fraction hired	Avg. exp. versus others' (1)	Store change (2)	"Endowed" workers (3)	"Exogenous" separation (4)
Black fraction quit/fired $\leq$ 3 months	-0.069	-0.067	-0.044	-0.039
	(0.023)	(0.045)	(0.027)	(0.018)
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	0.352	0.356	0.328	0.349
Standard deviation	0.266	0.309	0.172	0.098
Observations	30,985	977	11,659	34,025

NOTE. Clustered standard errors at the manager level are presented in parentheses. "Avg. exp. versus others'" includes a term for a manager's own experiences with black workers as well as terms for the experiences of other managers at the same store. "Store change" indicates that the manager changed store between their previous experience with black workers and the current hiring event. "Exogenous separation" restricts fires and quits to dissatisfaction with pay, compensation or benefits. "Endowed workers" corresponds to workers already in the department at the manager's arrival. See Table 2 for additional details.

resulted from the manager's behavior towards workers after hiring also decreases subsequent hiring, looking at specific reasons for separation. We estimate equation (2) restricting our measure of negative experiences to dismissals due to dissatisfaction with pay, compensation, or benefits, which are not controlled by the department manager. We find that a negative experience with black workers in the previous hiring event decreases black hiring in the current event, consistent with decreases being driven by managers updating their perceptions of worker groups rather than reflecting their own behavior or 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, though the time span of our

data is limited (roughly 6.5 years). This sharp degree of specificity largely reflects evidence on experience effects in other contexts (Malmendier, 2021a,b).

5 Non-experiential explanations

Our main proposition is that managers' beliefs about the performance of worker groups are shaped on the job. We have found that hiring experiences with black workers impact a manager's subsequent hiring, and that negative experiences in particular persistently reduce black hiring. Moreover, we have documented that managers' early hiring experiences yield the greatest impact on their subsequent hiring decisions, and that they seemingly adjust their hiring threshold for workers of different races based on their experiences.

When evaluating alternative explanations, one set of concerns regards factors correlated with the manager, particularly the manager's team. For instance, if workers refer 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 persistence in the race of a manager's hires that is not based on any factor specific to the manager.

Our strategy for evaluating these alternatives 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.

In particular, our results are inconsistent with any mechanism that operates through workers and customers. Table 4 highlights that the impact of negative experiences with black workers in a given department varies specifically with the timing of a manager's experiences. The empirical setup is analogous to 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. Other time placebos find little relationship between the hiring and performance of black workers in a department before a manager begins in their position and the manager's own hiring (Table C6). Similarly, there is a negligible relationship between the existing fraction of black workers in a department and the fraction of black workers hired in

a given event (Table C7). That is, a higher fraction of black workers, by itself, is not associated with a subsequent increase in black hiring by a manager when controlling for factors like the store location.

Moreover, several non-experiential mechanisms predict that hiring outcomes 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).[11] Yet, Table 5 shows the opposite, 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 than department level.

Second, we examine the possibility that the effects we identify are specific to the hiring manager, but are time invariant, in that they do not arise from them updating beliefs about the performance of worker groups from experience. In the presence of prejudice, negative (positive) experiences could reflect a bad (good) working climate for minority workers which translates to less (more) hiring. Even if we are interested in the manager's perception 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 bias against a group could affect both their expected tenure, for example by affecting how they are evaluated by the manager, and the likelihood of the manager hiring from the group.

Several of our results are inconsistent with static pre-existing biases or prejudice. First, the share of black hires by a manager in their first or previous hiring event does not help predict the probability of having negative experiences with subsequent black hires, consistent with hiring responding to experience rather than reflecting a fundamental bias (Table C6). Second, if hiring outcomes are driven by underlying

[11] 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.

bias which correlates with subsequent hiring, both negative and positive experiences should have a similar persistence, contrary to our results. Negative and positive experiences should also be positively serially correlated, inconsistent with Table 5. Third, these alternatives provide little rationale for the decreasing impact of hiring experiences over a manager's career or the fading impacts of experiences over subsequent hiring events shown in Figures 2 and 4. Fourth, since the rate of negative experiences with black and white workers is fairly similar, the key difference appears to lie in how managers respond to their experiences with these groups. Fifth, 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 (Table 6).[12] Sixth, several of our specifications include manager fixed effects to account for time-invariant differences in group hiring across managers.

To be clear, we do not interpret our results as an indication that managers may not be biased against black workers in ways beyond that which arise through updating from experience. In fact, insofar as other sources of bias such as those documented in Glover et al. (2017) and Sarsons (2019) arise from previous interactions, 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 group biases, as considered in much of the discrimination literature, are inconsistent with the evolution of hiring patterns across managers that we observe.

We may worry that our results could be driven by workers and managers sorting. Yet, much of the previous reasoning applied against other alternatives applies to workers selectively applying for positions with managers based on their history. In addition, negative experiences with black workers inherited, rather than hired, by the manager also decrease subsequent black hiring (Table 6). More fundamentally, 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

[12] Along with 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.

interviews, also suggest that workers are not assigned to a department manager based on their previous experiences with worker groups, especially since we exclude transfers.

Lastly, we may be concerned that demand shocks affect both hiring and retention efforts, and retention efforts affect our measure of positive or negative experience. To be a concern, these demand shocks would need to be race-specific, or at least correlated with race, such that both hiring and retention intensity for white or black workers grows for a race in a way that is not captured by store or time effects. For instance, suppose that high unemployment rates prompt managers to work less hard at retaining black workers and also make them less likely to hire new black workers; then managers will appear to have more negative experiences with black workers before hiring less of them. However, we find no substantial correlation between the likelihood of a negative experience and the unemployment rate, nor evidence that the effect of negative experiences on future hiring within a group depends on the unemployment rate (Table C8, both sets of estimates are also of the opposite sign than predicted by this alternative explanation). Moreover, if department and race-specific demand shocks were driving the observed correlation between a race's retention and hiring, then we would be able to reproduce our results in the placebo analyses presented above regarding the timing of a manager's arrival and their experiences. In contrast, our results are best explained by a manager's tenure and the nature of their idiosyncratic experiences, rather than outside factors.

To summarize, the explanation that best jointly rationalizes our results is managers updating their beliefs about the performance of black workers through their own hiring experiences with these workers. This updating 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. In the next section, we show that if managers attribute some of the discrepancy between a worker's expected and realized productivity to potential differences between worker groups, then a simple theoretical framework of belief updating is sufficient to explain the hiring patterns we document.

6 Theoretical framework

To provide a theoretical grounding for our findings, we propose a framework adapted from Lepage (2022) that incorporates experience effects in hiring into a statistical discrimination framework. The chief ingredients of the framework are that 1) managers face initial uncertainty about the productivity distribution of worker groups, 2) face greater uncertainty about the productivity of minority groups, and 3) update their beliefs based on their hiring experiences. Notably, unlike canonical theories of taste-based and statistical discrimination, the model provides an explanation for why hiring evolves with personal experience, and further offers an explanation for why the reliance on personal experiences systematically decreases minority hiring.

Consider a manager tasked with hiring the most productive workers available, taking vacancies, entry wages, and applicant pools for each position as given (as department managers at the firm do). They hire from an applicant pool of two groups denoted by W and B, B being a minority group. The expected productivity of worker i from group g, x_{ig}, depends on a noisy signal of individual productivity s_{ig} observed prior to hiring and group membership. The individual signal is composed of the worker's productivity and an unbiased noise component: $s_{ig} = x_{ig} + \varepsilon_{ig}$ with $\varepsilon_{ig} \sim N(0, \sigma^2_{\varepsilon g})$. For example, it could include information from a resume, pre-employment test, or interview.

Worker productivity is normally distributed with mean μ_g and variance σ^2_g, such that $x_{ig} \sim N(\mu_g, \sigma^2_g)$. Managers are initially uncertain about the productivity distribution of groups and have particularly weak priors about the productivity of minority groups.[13] In particular, to simplify exposition, we assume that managers know the productivity distribution of group W, the variance of group B productivity σ^2_B, and the noisiness of individual signals $\sigma^2_{\varepsilon g}$ for both groups.[14] This allows us

[13] Noisier information about the productivity of minority workers is consistent with work in statistical discrimination (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009). The main distinction is that we do not focus on whether individual signals from minority workers are noisier. Rather, we focus on managers facing uncertainty about the underlying productivity distribution of minority workers and its dynamic implications for subsequent hiring and belief updating. This uncertainty could arise because managers less frequently encounter them in the labor force and so have fewer experiences to inform their beliefs. Alternatively, it could arise from in-group/out-group dynamics, leading managers to have better initial information on their own group.

[14] Employers learning about productivity variance or individual signal precision affects the weight attached to individual signals versus group membership but leaves substantive implications unchanged.

to abstract from learning about group W to focus on manager beliefs about the mean productivity of group B, $\hat{\mu}_B$, which we posit is the key driver of hiring discrimination. To be clear, unlike typical statistical discrimination frameworks, we posit that managers are not only uncertain about the productivity of an individual group B worker, x_{iB}, but also about the mean of the distribution from which x_{iB} is drawn. Managers have prior beliefs about μ_B when they enter the labor market, these may be incorrect, and managers update their beliefs $\hat{\mu}_B$ based on the productivity of their hires.

Define $S_{mh} = \{x_{iBn} : i \text{ from } B \text{ is hired by } m \text{ to fill vacancy } n\}_{n=1}^{h}$ as the information set about workers from group B available to manager m after h hires. The expected productivity of worker i from group B is

$$P_{iBmh} = E[x_{iB}|s_{iB}, E[\hat{\mu}_B|S_{m,h-1}]] = \gamma_{Bmh}s_{iB} + (1-\gamma_{Bmh})E[\hat{\mu}_B|S_{m,h-1}]$$

where $\gamma_{Bmh} = \frac{\sigma_B^2 + Var[\hat{\mu}_B|S_{m,h-1}]}{\sigma_B^2 + Var[\hat{\mu}_B|S_{m,h-1}] + \sigma_{\varepsilon B}^2}$.[15] Managers maximize expected worker productivity over a total of H hiring decisions during their tenure as managers, $Max \sum_{h=1}^{H} \sum_{i=1}^{A_h} P_{igmh}$ where A_h is the set of applicants for vacancy h.[16]

To fill vacancy h, a manager hires the worker with the highest expected productivity out of applicant set A_h, with fraction F_{Bh} from group B. That is, worker i from group B is hired if $P_{iBmh} > P_{i'g'mh}$ for all $i' \in A_h$ from group g', and for $g' \in \{W, B\}$. Beliefs about group B's productivity carry over from the last hire when managers hire from group W, otherwise managers update their beliefs based on

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

[16] Since managers make repeated hiring decisions, they should value both the expected productivity of applicants as well as learning about group B productivity by hiring from the group, which is valuable because it can lead to better hiring decisions in the future. In our context, the value of learning itself is likely negligible: the median manager makes fewer than 14 hires during our sample period of roughly 6.5 years, expected tenure as a department manager is less than 10 years, and group B workers account for a minority of applicants such that they would be unlikely to fill most vacancies even if they were substantially more productive on average. We also find that managers do not respond differently to a first negative experience based on whether they will have to hire more workers in total over our sample period, inconsistent with behavior being substantively affected by internalized incentives to learn (Table C9). Early negative experiences with black or white workers also do not impact the total number of events that a manager will hire for over our sample period. In addition, previous work suggests that managers typically underestimate the value of learning when making hiring decisions (Li et al., 2020). Therefore, we focus on manager decisions informed by their beliefs about expected productivity, although the framework's broader takeaways remain largely unchanged when incorporating the value of learning (Lepage, 2022).

their group B hire's productivity. When updating their beliefs, managers first form an expectation about x_{iB} given that worker i was hired

$$E[x_{iB}|P_{iBmh} > P_{i'g'mh} \ \forall i' \in A_h, g' \in \{W, B\}].$$

Second, managers update their beliefs from $\hat{\mu}_B|S_{m,h-1}$ to $\hat{\mu}_B|S_{mh}$. The direction of updating depends on the discrepancy between the hire's expected and observed productivity

$$E[x_{iB}|P_{iBmh} > P_{i'g'mh} \ \forall i' \in A_h, g' \in \{W, B\}] - x_{iB}. \tag{3}$$

If realized productivity is above (below) expectation, denoted as a positive (negative) experience, managers update their beliefs upwards (downwards), increasing (decreasing) $E[\hat{\mu}_B]$ and therefore the probability that a group B worker is hired to fill subsequent vacancies. Importantly, the rate of hiring also drives the speed of learning. A positive experience, by increasing the probability of subsequent hiring from the group, increases the probability of observing signals about the group's productivity which leads to more accurate beliefs. In contrast, a negative experience decreases the probability of observing subsequent signals, leading to more persistent impacts on beliefs. Lastly, since $A_{Bh} < A_{Wh}$, group B workers are more infrequently hired even in the absence of bias, amplifying the persistent impact of negative experiences relative to positive ones since it is harder (easier) for managers to seek out (avoid) group B following positive (negative) experiences.

To summarize, managers update their beliefs about the performance of worker groups based on the observed productivity of their hires, and in turn hire based on these beliefs. Under the condition that belief updating is reflected through hiring decisions, this framework predicts the three main hiring patterns that we document: 1) positive/negative experiences of a manager with black workers positively/negatively affects their share of future black hires, 2) experiences of a manager with black workers disproportionately affect their future hiring compared to those with white workers, and 3) negative experiences of a manager with black workers have a more persistent impact on their share of future black hires than positive ones or experiences with white workers. In addition, the framework is consistent with our additional results: 1) belief updating becomes weaker as a manager's experience hiring workers of a given race grows and 2) the threshold value of the individual signal

s_{iB} that a manager requires to hire a black worker evolves with experience, creating selection in the productivity of subsequent black hires.

6.1 Discussion

The framework does not rely on prior bias, prejudice, true group differences in productivity, or deviations from Bayesian updating over a manager's own experiences. In fact, our empirical evidence suggests that managers who hire black workers had roughly unbiased priors about their productivity. If they systematically underestimated it, then negative experiences may have had a more muted impact on subsequent hiring and positive experiences may have lead to persistent increases.

Still, other biases in hiring and belief updating could also be at play. For example, greater updating by managers following their experiences with black workers is also consistent with stereotype formation (Allport et al., 1954). Psychologists assert that whiteness is largely invisible; Blacks are judged together while whites are assessed as individuals or along nonracial categories (Macrae and Bodenhausen, 2000). Since managers seemingly update more following both negative and positive experiences with black workers, our results do not appear driven by a static updating bias (e.g. as in Sarsons, 2019), although we cannot rule out that such biases also affect updating. Some of the dynamic hiring patterns we document are also consistent with biases relating to reinforcement learning with memory, selective recall, or motivated beliefs (Kunda, 1990; Bordalo et al., 2023). For example, a manager's experiences with minority workers, especially bad experiences, could be salient in their memory and affect hiring for a certain period of time before decaying. Alternatively, managers could interpret their experiences through an underlying bias or preference that makes them discard positive experiences with minority workers faster than negative ones.

We are limited in our ability to tease out the presence of these potential biases in the way managers update their beliefs or assess workers. Therefore, our preferred approach is to present a framework with the minimal number of ingredients necessary to explain our findings while adding to the classical statistical discrimination framework. More importantly, the key novel takeaway of our analysis remains across these different potential sources of bias: hiring discrimination evolves based on individual employer experiences on the job.

7 Conclusion

Our analysis is motivated by the basic question of whether discriminatory beliefs about minority groups among hiring managers are formed by the time they reach their position of hiring authority, or whether individual variation in the race of their hires results from their individual experiences on the job. Using data on a major retailer, we find evidence of the latter: whether experiences hiring minority workers are positive or negative seeds the race of future hires. Results are most pronounced for managers' initial black hires. When a black worker hired for a permanent position is fired or quits within three months, the likelihood that the manager hires another black worker drops substantially and persistently. Positive hiring experiences with black workers, in contrast, increase black hiring though the effects are relatively short lived. Results for white workers tend to mirror those for black workers, but are far weaker in magnitude and persistence.

Although our study focuses on one firm, the mechanism we document appears to be a product of managers having broad hiring authority within a screening process typical among large organizations. Our results suggest that delegating such authority to individual managers is not only prone to bias, but also inefficiency, as managers draw from relatively little experience and information when making hiring decisions. Moreover, the firm's organization and the labor market in general appear to provide little corrective information to managers with individual idiosyncrasies in their minority hiring fueled by personal experience. As a result, policies aimed at centralizing hiring, aggregating learning as is implicitly done with pre-employment testing and algorithmic hiring, or encouraging minority hiring through policies like affirmative action, may both enhance efficiency and reduce bias.

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.

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.

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

Bergemann, Dirk and Juuso Valimaki (2006) "Bandit problems," Cowles Foundation Discussion Paper No. 1551.

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, Peter Hull, and Alex Imas (2022) "Systemic discrimination: Theory and measurement,"Technical report, 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.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer (2023) "Memory and probability," *The Quarterly Journal of Economics*, 138 (1), 265–311.

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.

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.

Darity, William A and Patrick L Mason (1998) "Evidence on discrimination in employment: Codes of color, codes of gender," *Journal of Economic Perspectives*, 12 (2), 63–90.

Denrell, Jerker and James G March (2001) "Adaptation as information restriction: The hot stove effect," *Organization science*, 12 (5), 523–538.

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.

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.

Kline, Patrick, Evan K Rose, and Christopher R Walters (2022) "Systemic discrimination among large US employers," *The Quarterly Journal of Economics*, 137 (4), 1963–2036.

Kunda, Ziva (1990) "The case for motivated reasoning.," *Psychological bulletin*, 108 (3), 480.

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.

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 (2022) "Experience-based 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.

Li, Danielle, Lindsey Raymond, and Peter Bergman (2020) "Hiring as Exploration,"Technical Report 3630630, National Bureau of Economic Research.

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

Macrae, C Neil and Galen V Bodenhausen (2000) "Social cognition: Thinking categorically about others," *Annual review of psychology*, 51 (1), 93–120.

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*, 19 (6), 2857—-2894.

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

March, James G (1991) "Exploration and exploitation in organizational learning," *Organization science*, 2 (1), 71–87.

Miller, Conrad (2017) "The persistent effect of temporary affirmative action," *American Economic Journal: Applied Economics*, 9 (3), 152–90.

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.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz (2012) "The short-and long-term career effects of graduating in a recession," *American Economic Journal: Applied Economics*, 4 (1), 1–29.

Pager, Devah and Diana Karafin (2009) "Bayesian bigot? Statistical discrimination, stereotypes, and employer decision making," *The Annals of the American Academy of Political and Social Science*, 621 (1), 70–93.

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.

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.

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.

Online appendices for

"Learning to Discriminate on the Job"

Appendix A Worker performance and tenure

Our main results use short and long tenures to measure negative and positive experiences hiring workers of different races. Our use of turnover is based on the proposition that departures can be used as "revealed preference" for whether the manager and worker would like to continue the employment relationship. In this way, our approach follows canonical models that imply that the duration of an employment spell is governed by the match quality.

Our data also permit us to validate our turnover-based approach with other markers of whether the manager had a positive or negative hiring experience, which we explore in Table A1. First, for a subset of 7,606 commissioned salespeople hired during our sample period, we observe the monthly revenue associated with their sales divided by their sales targets, which is the main performance measure for these positions. To account for skewness and leverage, we winsorize this variable and take its logarithm. For the 102,746 person-month observations among salespeople, we regress sales performance as a function of our tenure-based experience measures: whether the worker was quit of fired within 3 months or achieved at least one year of tenure in their position, as well as location and month fixed effects. Workers who achieved tenure of at least 12 months had about 11 percent higher monthly sales versus their target compared to workers who were fired or quit within three months; not only did these hires last longer, they were more likely to outperform their targets while employed. This outcome conforms to the standard proposition that higher match qualities will beget both longer job tenures and greater productivity, both desirable outcomes for the hiring manager.

We also observe the HR manager's reported reason for which the worker turned over. For this variable, turnover is categorized as voluntary or involuntary, and these broad categories include detailed subsets of classifications Although the distinction between types of turnover can be subjective, our approach relies on the assumption that involuntary turnover generally suggests a more negative experience, particularly terminations for poor performance. For these analyses, we return to the full sample

of jobs considered in the main analysis, but restrict the sample to 240,176 workers who left within 3 months or after 12 months. Among observed departures, we then examine the reasons for turnover as outcomes. We find that 22.8% of workers who left within 3 months were involuntarily terminated, versus 16.2% who those who achieved at least 12 months of tenure. More strikingly, workers who depart within 3 months were twice as likely to be terminated for poor performance. These statistics suggest short tenures are more likely to reflect a negative experience for the manager.

Finally, we evaluate separations classified as an employee-initiated voluntary departure for better opportunities under the premise that such employee-initiated separations are more likely to correspond to a positive experience for the period that the worker was employed. Workers who leave within 3 months are substantially less likely to be classified as quits for better opportunities compared to those who had tenures of at least 12 months.

Put together, evidence from the sales and turnover data corroborate that short and long tenures can be used as a measure of negative and positive experiences.

Table A1: Worker performance and experience measures

Related performance outcomes	Mean	Std. error	P-value
(a) Sales performance relative to target			
Among workers who quit or were fired within 3 months	-0.136	0.017	< 0.0001
Among workers who achieved tenure of at least 12 months	0.005	0.017	
(b) Sales performance relative to target, first 3 months			
Among workers who quit or were fired within 3 months	-0.136	0.017	< 0.0001
Among workers who achieved tenure of at least 12 months	0.027	0.018	
(c) Worker was involuntarily terminated			
Among workers who quit or were fired within 3 months	0.228	0.001	< 0.0001
Among workers who achieved tenure of at least 12 months	0.162	0.001	
(d) Worker was terminated for unsatisfactory performance			
Among workers who quit or were fired within 3 months	0.116	0.001	< 0.0001
Among workers who achieved tenure of at least 12 months	0.058	0.001	
(e) Worker quit for better opportunities			
Among workers who quit or were fired within 3 months	0.068	0.001	< 0.0001
Among workers who achieved tenure of at least 12 months	0.169	0.002	

NOTE. This table shows alternative performance measures and their relation to our two main measures of good and bad experiences. Outcome (a) restricts the data to commissioned salespeople who were hired and either fired or quit within 3 months or achieved at least 12 months of tenure in the position for which they were hired. Each worker's sales performance is calculated monthly, corresponding to 102,746 person-month observations. The measure corresponds to the log of monthly commissioned sales in dollars divided by the worker's sales target and is also purged of store and month effects. A value of zero means that the worker exactly hit their target. Outcome (b) is similar to outcome (a) but restricts to performance during the first three months of employment even for workers who eventually achieve tenure of at least 12 months. Outcomes (c), (d), and (e) use all jobs represented in our main analysis, but restrict the sample to 240,176 individuals who were hired and terminated during our sample period, either because they were fired or quit within 3 months or because they separated from their position after achieving at least 12 months of tenure. Turnover reasons are reported by HR representatives, and include involuntary reasons and voluntary reasons; outcome (c) includes all involuntary reasons, whereas outcome (d) focuses on a subset of involuntary terminations relating to unsatisfactory performance and attendance. Outcome (e) includes voluntary separations for the worker's career advancement or return to studies. Robust standard errors are presented in the second column. P-values are for the test that the two adjoining means are equal.

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 vary by store, department, job, time period, or economic condition, we take Abowd et al. (1999)'s approach of analyzing connected sets of workers.[17] 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

$$Black_{imjlt} = X_{mjlt}\beta + \gamma_m + \alpha_j + \lambda_l + \theta_t + \varepsilon_{imjlt} \tag{4}$$

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.[18] 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.

Figure B1 presents the raw shares, while B2 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,

[17] 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)

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

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 identity of the hiring manager is an important predictor of black hiring in a department. Figure B3 presents analogous results for white hiring while Figure B4 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 B5. 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 B6, 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 the shares of black and white hires, by manager

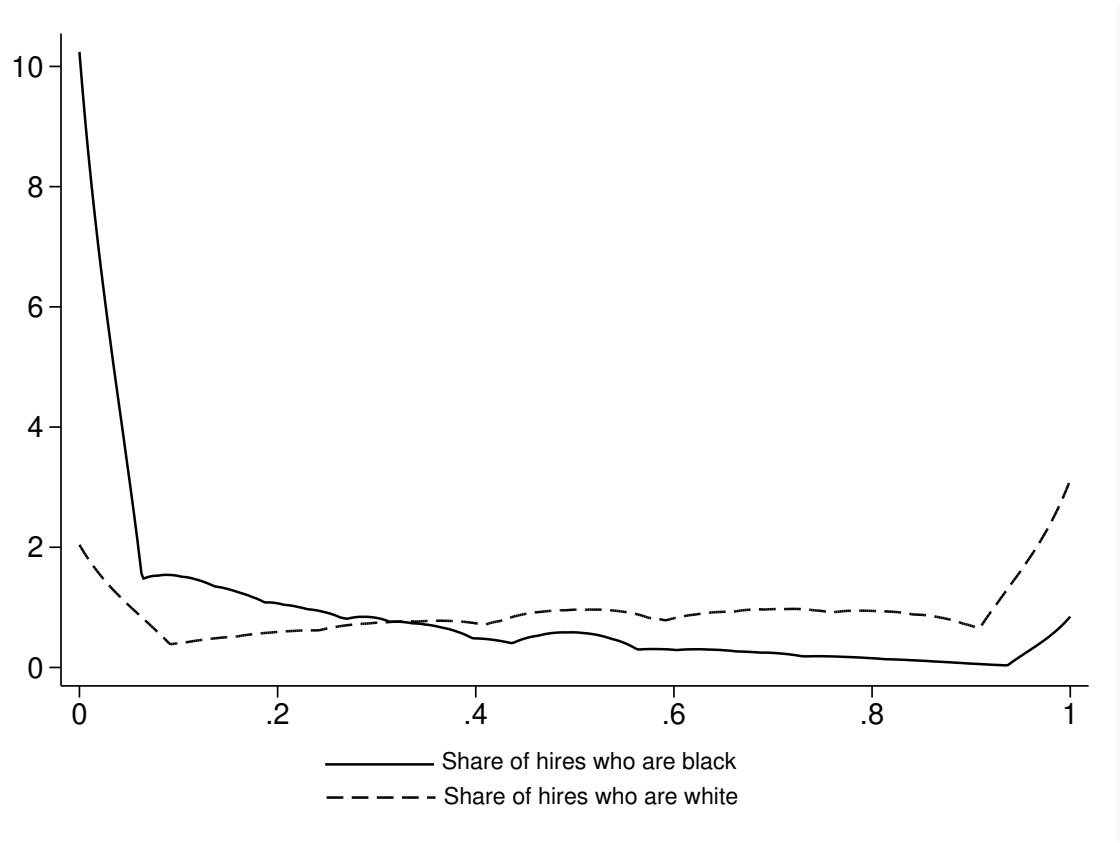

NOTE. Black share bandwidth: 0.066. White share bandwidth: 0.091.

Figure B2: Kernel density estimates of predicted black hiring shares

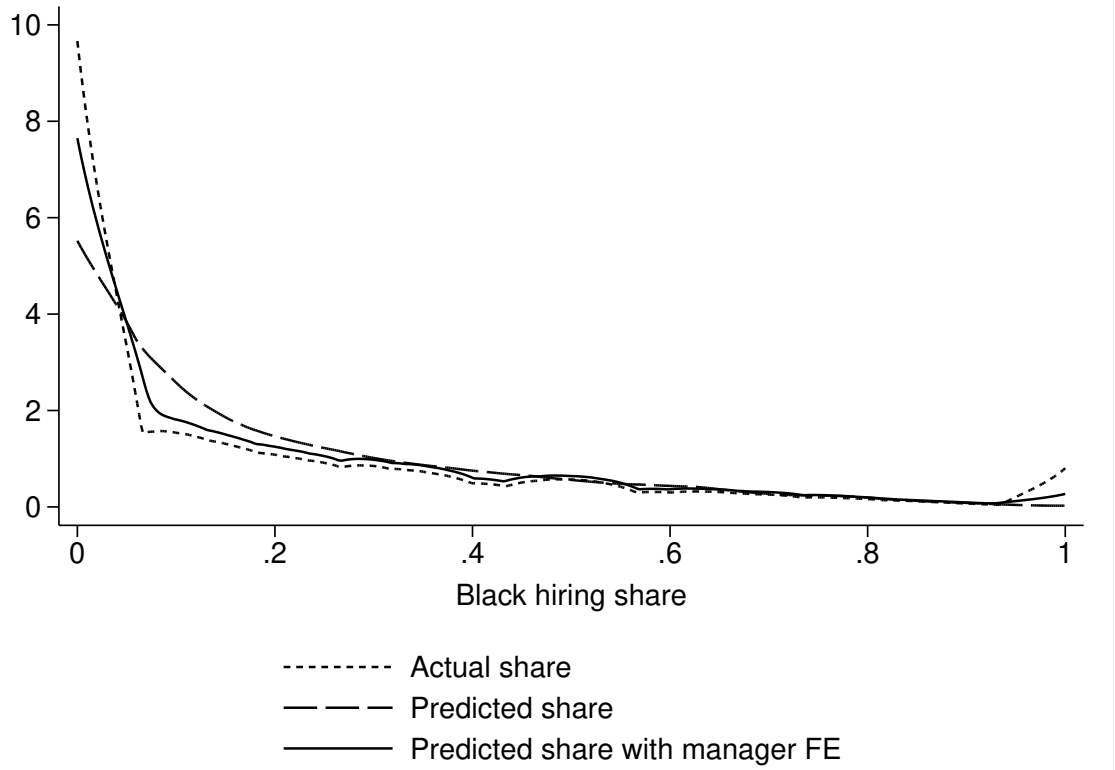

NOTE. Predicted shares are obtained by averaging predicted values for each manager from an individual hire level linear probability model regression including whether the worker was hired for a part-time or full-time job, the manager's previous number of hires at the time that the current worker is hired, 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 B3: Kernel density estimates of manager predicted white hiring shares

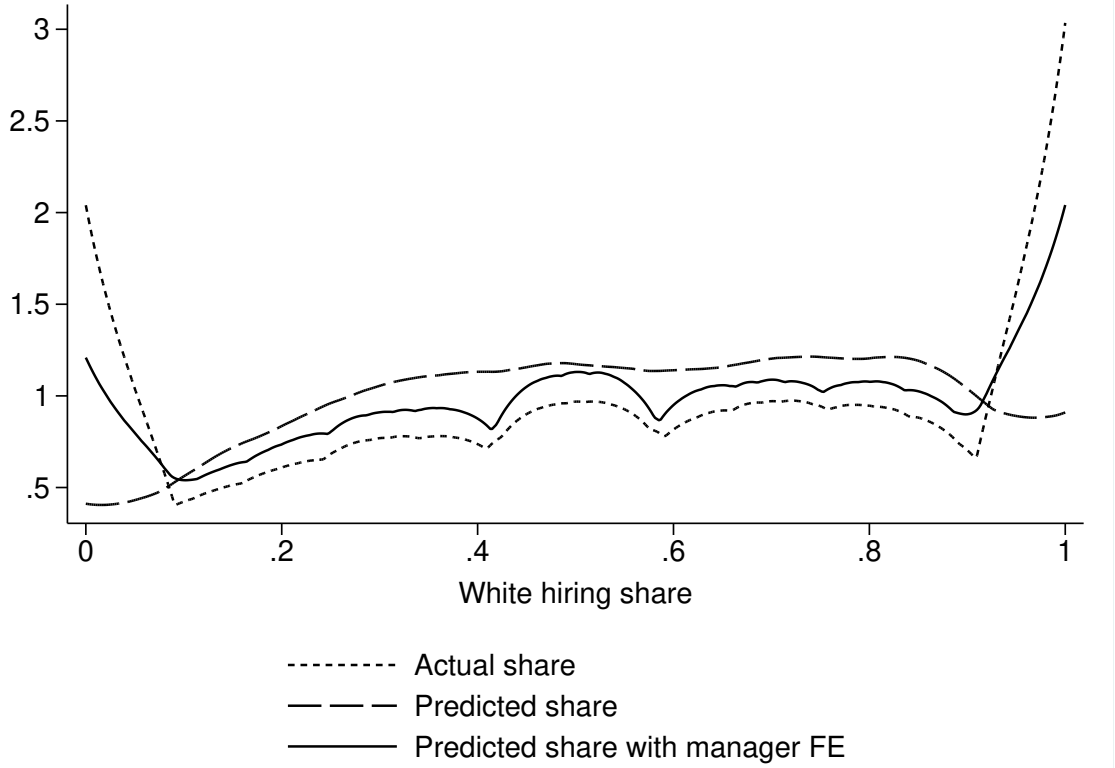

NOTE. See Figure 2 for details. Actual share bandwidth: 0.091. Predicted share bandwidth: 0.074. Predicted share with manager FE bandwidth: 0.086.

Figure B4: Kernel density estimates of manager predicted black hiring shares, managers with over 5 hires

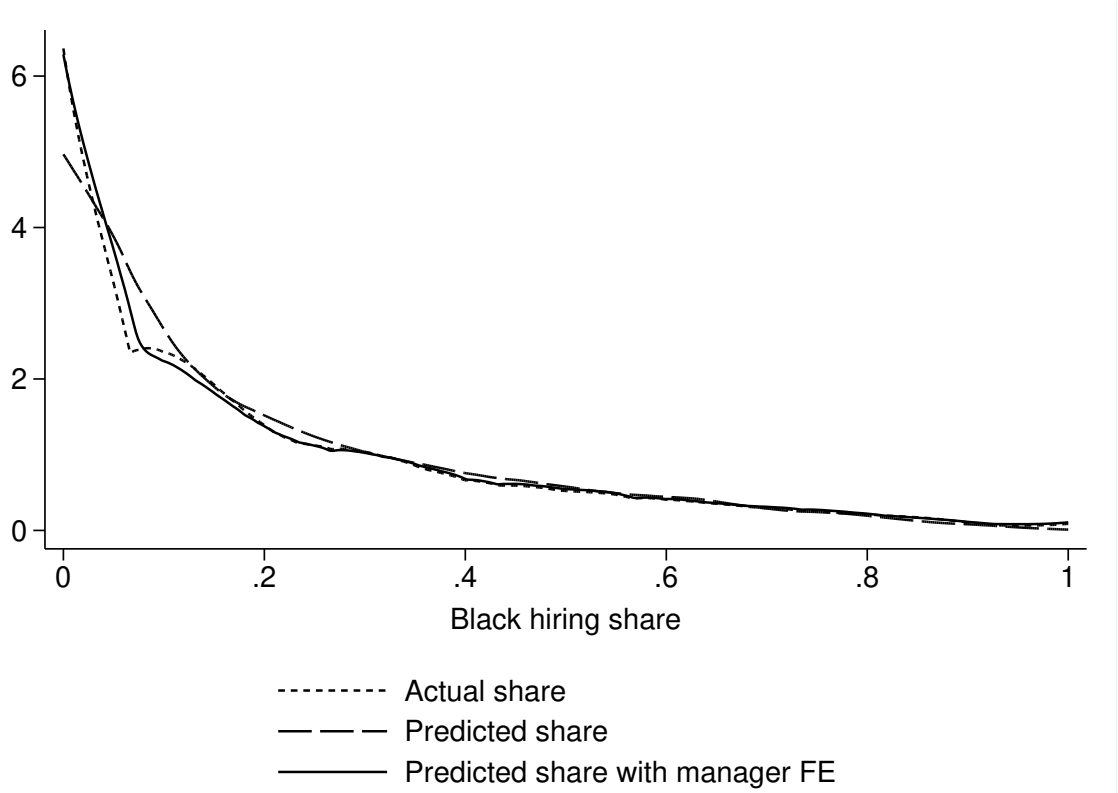

NOTE. See Figure 2 for details. Actual share bandwidth: 0.066. Predicted share bandwidth: 0.063. Predicted share with manager FE bandwidth: 0.068.

Figure B5: Distribution of manager fixed effects for black hiring

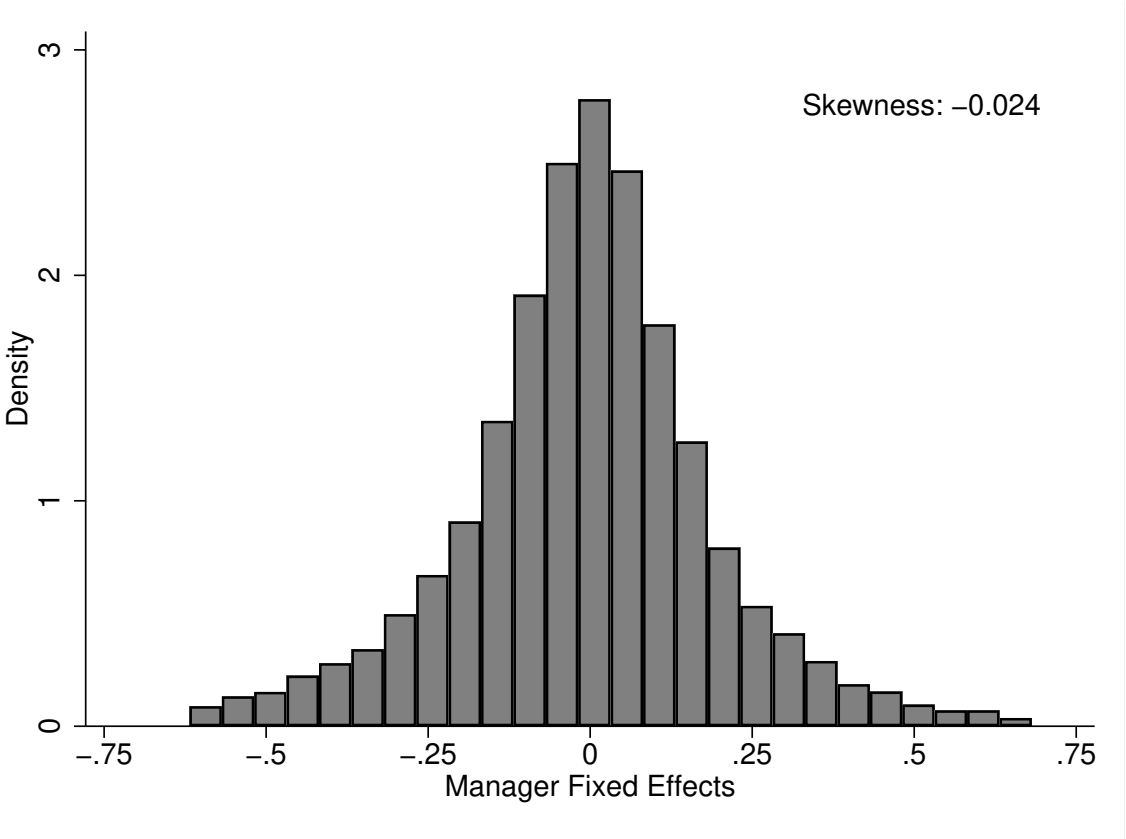

NOTE. See Figure 2 for specification details. Fixed effects are estimated for the largest connected sample of stores and managers following Abowd et al. (1999) and adjusted using empirical Bayes shrinkage.

Figure B6: Distribution of manager fixed effects for white hiring

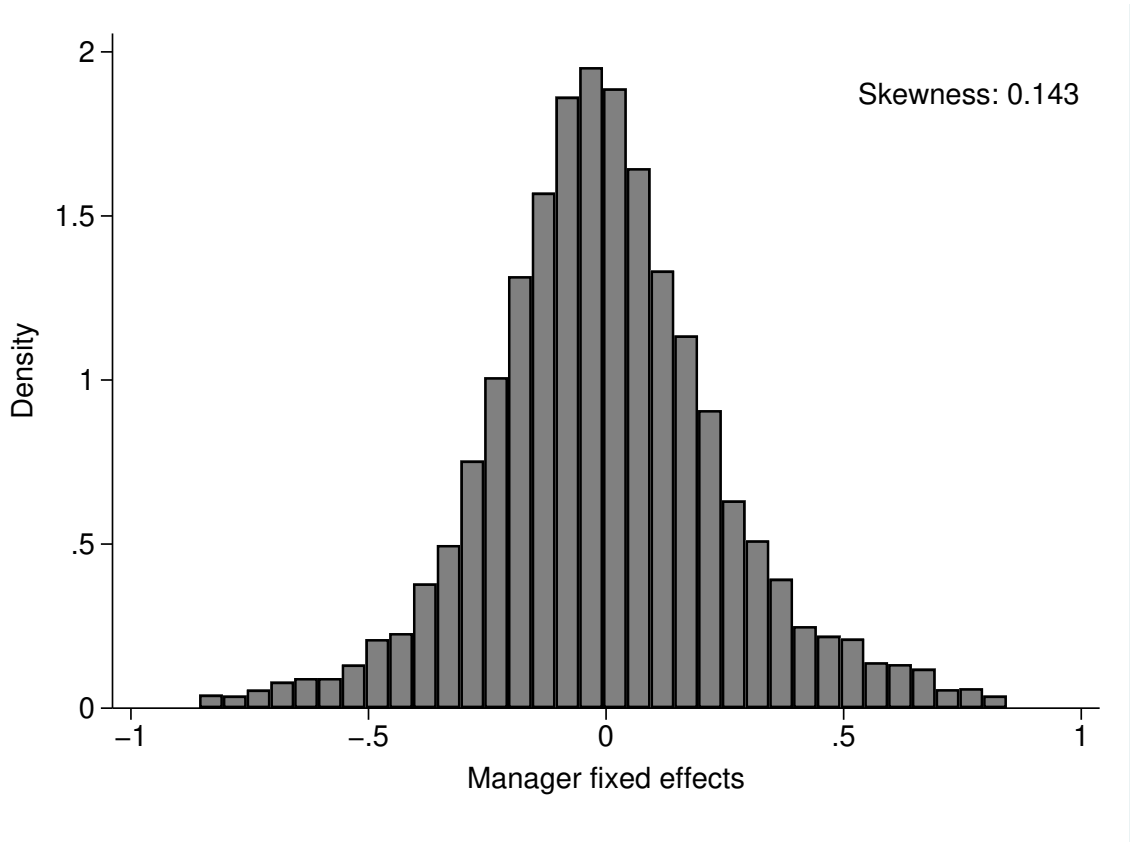

NOTE. See Figure 2 for details.

Appendix C Additional results on the impact of hiring experiences

Figure C1: Tenure of first hire across managers

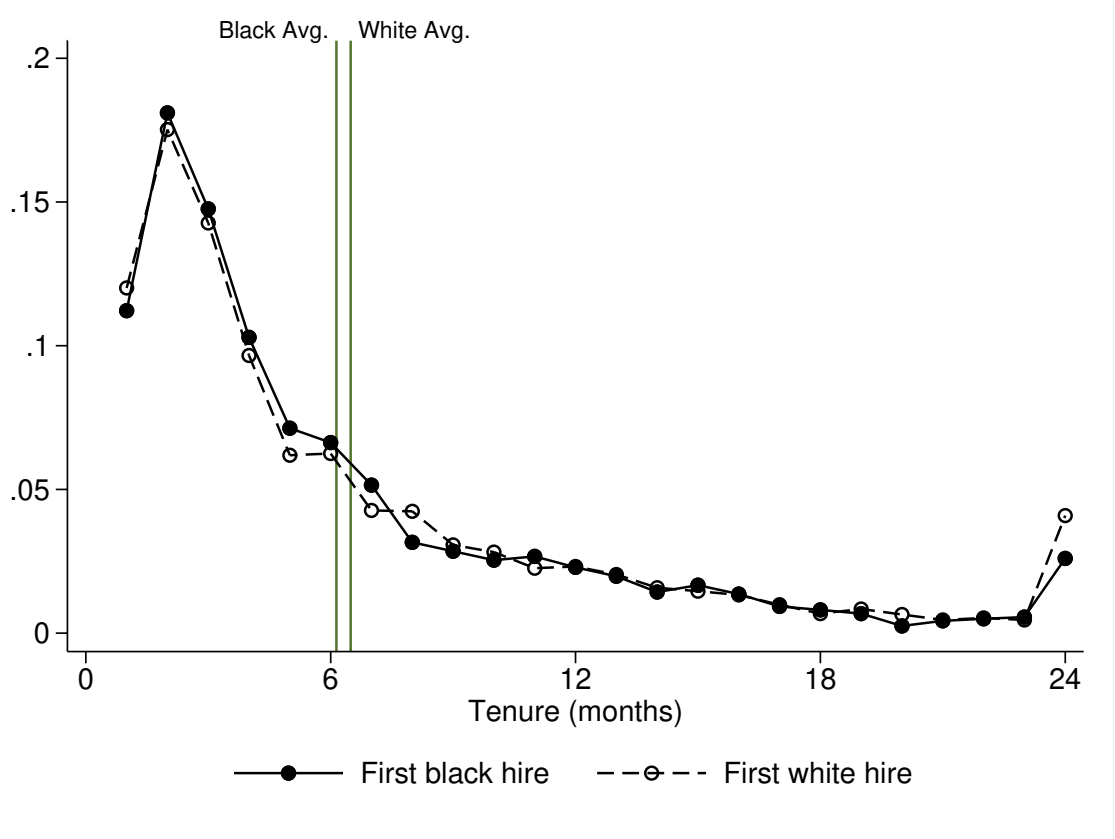

NOTE. Workers hired in the last two years of our sample are excluded since they cannot achieve the maximum tenure censored at 24 months. Hiring events are restricted to those with at most one worker hired from each racial group.

Figure C2: Impact of a first negative hiring experience with black workers on black hiring share, by local black population

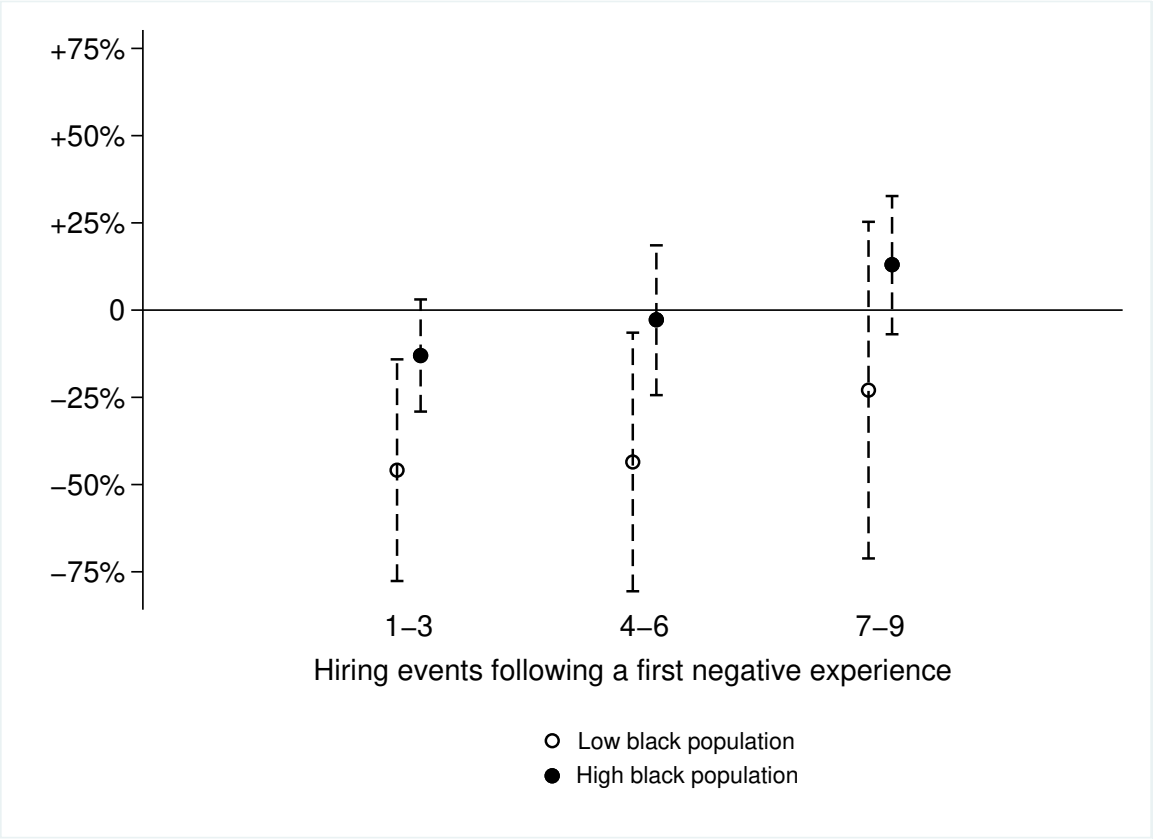

NOTE. 95% confidence intervals from clustered standard errors at the manager level are presented using dashed lines. High and low black population areas correspond to stores in ZIP codes above and below the median ratio of black-to-white population. Estimated coefficients and confidence intervals are normalized by the average black hiring share in low and high black population areas to obtain an effect in percentage accounting for the large baseline differentials in black hiring across areas. 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. See Figure 2 for additional details. Source: U.S. Census Bureau, 2012-2016 American Community Survey 5-Year Estimates.

Table C1: OLS estimates of the cumulative impact of previous experiences with black and white workers on current black hiring, expected tenure

Black fraction hired	(1)
Black expected months of tenure	0.050
	(0.011)
White expected months of tenure	-0.009
	(0.011)
Manager FE	Y
Store FE	Y
Hiring month and year FE	Y
Worker and event controls	Y
Outcome mean	0.380
P-value: B = -1 * W	0.023
Standard deviation black	0.524
Standard deviation white	0.555
Observations	35,937

NOTE. Clustered standard errors at the manager level are presented in parentheses. Expected months of 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 C2: OLS estimates of the cumulative impact of previous experiences with black and white workers on current black hiring, bottom and top quartiles of tenure

Black fraction hired	(1)	(2)
Black expected tenure in the bottom quartile	-0.050	
	(0.007)	
White expected tenure in the bottom quartile	0.022	
	(0.007)	
Black expected tenure in the top quartile		0.065
		(0.007)
White expected tenure in the top quartile		0.026
		(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	0.380
Observations	35,883	35,883

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 C3: OLS estimates of the cumulative impact of previous experiences with black workers on current black hiring, 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 C4: OLS estimates of the impact of the latest experience with black and white workers on current black hiring, negative and positive experiences

Black fraction hired	(1)	(2)	(3)	(4)	(5)	(6)
Black fraction quit/fired $\leq$ 3 months	-0.041	-0.040	-0.041			
	(0.005)	(0.006)	(0.006)			
White fraction quit/fired $\leq$ 3 months			0.035			
			(0.006)			
Black fraction tenure $\geq$ 12 months				0.014	0.022	0.022
				(0.008)	(0.009)	(0.009)
White fraction tenure $\geq$ 12 months						-0.016
						(0.008)
Manager FE		Y	Y		Y	Y
Worker and event controls		Y	Y		Y	Y
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Outcome mean	0.367	0.367	0.348	0.366	0.366	0.347
Standard deviation black	0.409	0.409	0.409	0.296	0.296	0.296
Standard deviation white	0.352	0.352	0.352	0.297	0.297	0.297
Observations	34,496	33,971	31,911	29,511	29,064	27,249

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

Table C5: OLS estimates of the impact of the first three experiences with black and white workers on current black hiring, negative and positive experiences

Black fraction hired	(1)	(2)	(3)	(4)	(5)	(6)
Black fraction quit/fired $\leq$ 3 months	-0.043	-0.049	-0.042			
	(0.013)	0.013	(0.014)			
White fraction quit/fired $\leq$ 3 months			-0.013			
			(0.017)			
Black fraction tenure $\geq$ 12 months				0.035	0.009	0.020
				(0.018)	(0.017)	(0.018)
White fraction tenure $\geq$ 12 months						0.018
						(0.020)
Manager FE		Y	Y		Y	Y
Worker and event controls		Y	Y		Y	Y
Store FE	Y	Y	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y	Y	Y
Outcome mean	0.405	0.404	0.402	0.400	0.430	0.401
Standard deviation black	0.200	0.239	0.239	0.200	0.200	0.200
Standard deviation white	0.214	0.205	0.205	0.214	0.214	0.214
Observations	34,136	31,772	27,829	29,300	25,041	23,527

NOTE. Clustered standard errors at the manager level are presented in parentheses. Experience measures refer to the average hiring outcome over the first three hiring events with black and white workers (excluding managers who hired less than three times from either group). See Table 2 for additional details.

Table C6: Placebo tests

	Negative experiences before manager DV: Frac. hired (1)	Black hiring before manager DV: Neg. exp. (2)	Hiring in event t-1 and neg. exp. t DV: Neg. exp. (3)	Hiring in event 1 and neg. exp. t DV: Neg. exp. (4)
Black fraction quit/fired $\leq$ 3 months	0.018			
	(0.022)			
Black fraction hired		0.010	0.0001	-0.009
		(0.033)	(0.006)	(0.014)
Manager FE			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	0.478	0.219	0.265	0.225
Observations	9,741	7.827	36,256	19,200

NOTE. Robust standard errors are presented in parentheses for columns 1-2 and clustered at the manager level for column 3. The first column presents results of a time placebo test investigating the impact of a worse performance by black hires in the 3 months before a manager begins in their position at the department on the hiring of black workers by the manager in subsequent hiring events. The second column presents results of a time placebo test investigating the impact of more black hiring in the 3 months before a manager begins in their position at the department on the probability of a negative experience with black workers in subsequent hiring events. The third column presents results of a placebo test investigating whether the share of black hires in the previous hiring event predicts the probability of a negative experience with black workers in the current hiring event. The fourth column presents results of a placebo test investigating whether the share of black hires in the first hiring event predicts the probability of a negative experience with black workers in the current hiring event. See Table 2 for additional details.

Table C7: 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	-0.001
	(0.0004)
Manager FE	Y
Store FE	Y
Hiring month and year FE	Y
Outcome mean	0.286
Standard deviation black	0.254
Observations	99,581

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

Table C8: Placebo tests and unemployment rate

	DV: Black fraction hired		DV: Black fraction	DV: White fraction
	Low unemployment	High unemployment	quit/fired $\leq$ 3 months	quit/fired $\leq$ 3 months
	(1)	(2)	(3)	(4)
Latest black fraction quit/fired $\leq$ 3 months	-0.052	-0.036		
	(0.009)	(0.009)		
Latest white fraction quit/fired $\leq$ 3 months	0.033	0.037		
	(0.009)	(0.009)		
Unemployment rate			-0.886	-1.089
			(1.191)	(0.910)
Manager FE	Y	Y	Y	Y
Worker and event controls	Y	Y	Y	Y
Store FE	Y	Y	Y	Y
Hiring month and year FE	Y	Y	Y	Y
Outcome mean	0.342	0.355	0.266	0.208
Observations	16,478	15,110	35,879	34,731

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

Table C9: OLS estimates of the impact of the first experience on the manager's total hiring events and of the interaction between the first experience and the total hiring events on black hiring

	DV: Total Hiring Events	DV: Black fraction hired in the next 3 events
	(1)	(2)
Black fraction quit/fired $\leq$ 3 months	0.412	-0.021
	(0.373)	(0.007)
White fraction quit/fired $\leq$ 3 months	0.290	
	(0.363)	
Total hiring events		0.001
		(0.001)
Black frac. quit/fired $\leq$ 3 months x Tot. hiring events		-0.0001
		(0.001)
Worker and event controls	Y	Y
Store FE	Y	Y
Hiring month and year FE	Y	Y
Outcome mean	22.070	0.440
Observations	35,613	9,472

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.