Employer Learning, Biased Beliefs, and Labor Market Discrimination

by

Louis-Pierre Lepage

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2021

Doctoral Committee:

Professor John Bound, Chair
Professor Charles Brown
Associate Professor David Miller
Professor Tanya Rosenblat

Louis-Pierre Lepage

lepagelp@umich.edu

ORCID iD: 0000-0002-0626-4950

ACKNOWLEDGEMENTS

I would first like to thank my advisors John Bound, Charlie Brown, David Miller, Tanya Rosenblat, and Jeff Smith for their invaluable feedback as well as their continued guidance and support. I would also like to thank my co-authors Alan Benson and Steve Lehrer for their advice and wonderful collaboration as well as several other faculty members at the University of Michigan who have helped me.

I graciously acknowledge financial support from the University of Michigan's Department of Economics, Rackham Graduate School, and Michigan Institute for Teaching and Research in Economics (MITRE), as well as the Social Sciences and Humanities Research Council (SSHRC).

I am thankful to my parents André and France as well as my grandmother Claire for their constant support and encouragement throughout my studies. Lastly, I am thankful to my fiancée Rosina for her great professional support and her unwavering personal support.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

My dissertation proposes and tests a new theory of labor market discrimination based on employers developing persistent negatively biased beliefs about the productivity of worker groups through their hiring experiences with these groups.

My first chapter presents a statistical discrimination model in which employers are initially uncertain about the productivity of groups and endogenously learn about it through their hiring experiences. An employer's hiring history determines their beliefs about group productivity, but also shapes their subsequent incentives to hire from the group and learn more about their productivity. Positive experiences create positive biases which correct themselves by leading to more hiring and learning. Negative experiences create negative biases which decrease hiring and therefore learning, leading to the persistence of negative biases. Differential hiring and learning across employers thus generates a negatively-skewed belief distribution about worker group productivity. Endogenous employer learning disproportionately affects workers from minority or underrepresented groups if there is less initial information available about their productivity in the labor market, making employers more reliant on their own experiences to assess these groups. I show that discrimination in the form of a wage below these groups' expected productivity can arise and persist from this initial information asymmetry even with market competition and without true productivity differences between groups, prior bias, or prejudice. The model generates analogous predictions to taste-based discrimination, in a statistical framework with beliefs replacing preferences, providing a new way to understand prejudice as the result of "incorrect" statistical discrimination. The model helps explain the persistence and

pervasiveness of discrimination, also generating new implications for policies like affirmative action which can induce employers to hire from specific groups and learn about their productivity.

My second chapter tests how hiring experiences of employers with worker groups impact hiring and beliefs about group productivity. I design an experiment where employers hire a worker from one of two groups each period, with one group framed as a minority about whose productivity employers are initially given less information. Employers are incentivized to hire productive workers, observe their hire's productivity after they perform a real-effort task, and then report their beliefs about group productivity. The results show that negative experiences with the minority group, captured through the hiring of low productivity workers, lead to negatively-biased beliefs about the group's productivity by decreasing subsequent hiring and learning. In contrast, positive biases which arise from positive experiences are mitigated through increased hiring, leading to a negatively-skewed belief distribution across employers.

My third chapter joint with Alan Benson at the University of Minnesota uses employment records of a large retail firm to study how hiring experiences of managers with worker groups influence their hiring. We study the hiring of black and white workers, relating current hiring decisions of a manager to measures of their previous experiences with these groups. We find that negative experiences with previous hires of a group, measured by a higher fraction quickly being fired or quitting, decrease subsequent hiring of the group. More positive experiences, measured by a higher fraction of previous hires achieving long tenure, increase subsequent hiring of the group. These impacts are substantively larger for black workers, and early negative experiences with them lead to particularly persistent decreases in relative hiring of the group.

CHAPTER I

Introduction

After decades of cultural change and anti-discrimination legislation, there remain substantial labor market outcome differentials across race and gender (Lang and Lehmann, 2012; Blau and Kahn, 2017). Models are key to understand the contribution of discrimination to these differentials. Statistical discrimination arises as a rational response to productivity differentials across worker groups, such that employers have correct equilibrium beliefs about group productivity (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977, Coate and Loury, 1993; Moro and Norman, 2004; Fang and Moro, 2011). Taste-based discrimination explains discrimination through exogenous preferences of employers for groups (Becker, 1957; Black, 1995), creating differences between average performance and average pay of groups. An alternative explanation is that discrimination arises from incorrect or biased beliefs of employers about the productivity of groups.[1] Then, discrimination does not reflect true group differentials, but arises from a lack of information or learning. The distinction between statistical discrimination with biased beliefs and other theories is critical, because they can lead to very different predictions regarding how discrimination arises and can be mitigated.

In this dissertation, I propose a new mechanism, endogenous employer learning, through which biased employer beliefs can arise and persist in the labor market and provide supporting evidence across two empirical settings. The fundamental idea is that if employers enter the labor market with uncertainty about the productivity of

[1]See Fershtman and Gneezy (2001), Reuben et al. (2014), Bordalo et al. (2016), Mobius et al. (2016), Laouénan and Rathelot (2017), Van Dalen and Henkens (2019), Arnold et al. (2018), Landsman (2018), Lesner (2018), Bohren et al. (2019a, 2019b), Bordalo et al. (2019), and Sarsons (2019).

different worker groups, then they may naturally update their beliefs as their hire and observe the productivity of individual workers from these groups. Then, hiring experiences with a group shape employer beliefs about their productivity, but also subsequent incentives of employers to hire from the group and learn more about their productivity. That is, an employers hiring and learning are endogenous to their previous hiring experiences with groups. Chapter II presents a new model of statistical discrimination showing that this can lead to persistent discrimination against worker groups whose productivity is initially more uncertain to employers. Chapter III presents supporting evidence on the endogenous formation of biased beliefs through the endogenous learning mechanism in an experimental labor market implemented through an online experiment. Chapter IV presents evidence that individual manager hiring decisions at a large US retail firm are influenced by their previous hiring experiences with worker groups in the specific ways predicted by the endogenous employer learning mechanism.

This dissertation presents and tests a new theory of discrimination with important implications for the theoretical literature, empirical work, and policy. First, it provides a new conceptual justification for the pervasiveness and persistence of biased employer beliefs in labor markets and documents empirical behavior consistent with such biased beliefs. Second, while biased beliefs in my framework arise from statistical discrimination by employers, their behavior does not reflect true worker group differences because of their biased beliefs, and resulting discrimination in fact has similar intuitive similarities with the alternative theory of taste-based discrimination. This new theory thus blurs the line between the two classical theories of discrimination in economics, which are often presented and understood as distinct and mutually exclusive. Third, since this type of discrimination arises from employers learning too little about the productivity of worker groups, it has distinct implications than previous theories regarding policies which can induce employers to learn more or provide them with additional information about worker groups.

CHAPTER II

Endogenous Learning, Persistent Employer Biases, and Discrimination

2.1 Introduction

After decades of cultural change and anti-discrimination legislation, there remain substantial labor market outcome differentials across race and gender (Lang and Lehmann, 2012; Blau and Kahn, 2017). Models are key to understand the contribution of discrimination to these differentials. Statistical discrimination arises as a rational response to productivity differentials across worker groups, such that employers have correct equilibrium beliefs about group productivity (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977, Coate and Loury, 1993; Moro and Norman, 2004). Taste-based discrimination explains discrimination through exogenous preferences of employers for groups (Becker, 1957), creating differences between average performance and average pay of groups. An alternative explanation is that discrimination arises from incorrect or biased beliefs of employers about the productivity of groups.[1] Then, discrimination does not reflect true group differentials, but arises from a lack of information or learning. The distinction between statistical discrimination with biased beliefs and other theories is critical, because they can lead to very different predictions regarding how discrimination arises and can be mitigated.

[1]See Fershtman and Gneezy (2001), Reuben et al. (2014), Bordalo et al. (2016), Mobius et al. (2016), Laouénan and Rathelot (2017), Van Dalen and Henkens (2019), Arnold et al. (2018), Landsman (2018), Lesner (2018), Bohren et al. (2019a, 2019b), Bordalo et al. (2019), and Sarsons (2019).

One potential source of biased beliefs in the labor market is initial employer uncertainty about the productivity of groups. That is, when employers enter the labor market, they are not only uncertain about the individual productivity of potential workers, but also the underlying productivity distribution of their group. If employers perceive that there may be a correlation between individual productivity and group membership, for example due to historical or social factors disadvantaging some groups, then they should value learning about group productivity since it can improve hiring decisions in a statistical discrimination framework. A natural source of employer learning about the productivity of groups is their own hiring experiences with workers of these groups. In this context, previous experiences of an employer with workers of a given group not only shape their beliefs about the group's productivity, but also their subsequent incentives to hire from the group and, indirectly, to learn more about their productivity. Moreover, employer learning about the productivity of minority or disadvantaged groups may be particularly important if there is less initial information available about them in the labor market, making employers more reliant on their own experiences to assess their productivity.

This paper presents a new model of statistical discrimination that captures these intuitive insights and highlights their implications for labor market discrimination. Employers perceive group membership as a potentially relevant indicator of individual productivity and endogenously learn about group productivity through their own hiring. In a dynamic setting, employers have noisier initial information on one group's productivity relative to another (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009) and trade off learning about that group against current-period profit maximization. A substantial component of the information observed through hiring is privately-observed by the hiring employer, such that an employer's hiring history shapes their future hiring and learning. Positive experiences create positive biases, which endogenously correct themselves through more hiring and learning. Negative experiences, however, create negative biases which decrease hiring and learning.[2] Differential learning across employers results in

[2]The dynamic decision problem I study has intuitive similarities with self-confirming equilibrium models for non-cooperative games (Fudenberg and Levine, 1993a; 1993b). Both study the outcome of a learning process in which agents learn from their experiences, beliefs are not contradicted along the equilibrium path, and inefficiencies arise from insufficient learning. My model focuses on learning about the environment rather than other players' strategies, showing that some employers optimally stop learning.

a negatively-skewed distribution of beliefs about the group's productivity.

Each period, employer beliefs determine market clearing wages, pinned down by the beliefs of the marginal employer, and optimal hiring therefore follows a cutoff rule in beliefs about a group's productivity. Employers below the cutoff do not hire from the group, preserving negative biases. The key prediction of the model is that, over time, the skewness in the belief distribution can cause the wage of the group about whose productivity employers have noisier initial information to fall and remain below their expected productivity in the long run. The model predicts discrimination due to uncertainty, even with equally productive worker groups and without prior biases or endogenous worker responses.[3] Further, since discrimination arises endogenously from profit-maximizing decisions, it can survive competition in the form of higher exit rates for biased employers if new entrants face a similar learning problem. In summary, heterogeneous biased beliefs persist within a statistical discrimination framework; they are not necessarily eliminated through learning or competition.

In the baseline model, employer learning is private and employers do not learn about group productivity from observing the evolution of wages. Labor markets may provide few salient signals to an employer who decides to stop hiring from a group based on their own experiences. First, evidence indicates that employers weight their own experiences particularly heavily, because hiring outcomes of other employers are partially unobservable, employment contexts differ across employers, and employers overestimate the precision of their own information (Waldman, 1984; Moore et al., 2015; Ge et al., 2020; Guenzel and Malmendier, 2020). Second, relative wages in practice summarize how the market clears from mostly unobserved decentralized bilateral bargaining outcomes, rather than an aggregate price signal, and depend on many factors beyond employer beliefs like sectoral and occupational shifts, macroeconomic shocks, and demographic changes. Accordingly, employers routinely isolating the residual wage component that is due to changing subjective employer beliefs about group productivity appears unrealistic, as exemplified by the decades-long debate

[3] Arrow (1973) mentions that biased priors could lead to a self-fulfilling prophecy if employers ignore subsequent information or worker responses confirm employer beliefs, but these models have no learning. I provide a mechanism through which biased beliefs create discrimination without biased priors, deviation from profit-maximization, or endogenous worker investments.

on wage decomposition in economics (Lang and Lehmann, 2012).[4] Nevertheless, I consider an extension in which employers noisily learn from outside sources such as other employers or wages. Discrimination can still persist if employers put more weight on their own experiences and there is either dynamic entry of employers or evolving productivity of worker groups. Accordingly, an intuitive interpretation of the model is a cohort of employers learning about a cohort of workers, with imperfect transfer across cohorts.

Unlike classical statistical discrimination models, I do not assume that employers know group productivity or have correct equilibrium beliefs about it. Rather, I model learning about groups, resulting in heterogeneous negatively biased beliefs which arise from uncertainty about the information environment and can persist with market clearing and endogenous wages. My model highlights that learning about some groups can be slow, complementing the employer learning literature which focuses on learning about individuals within groups (Farber and Gibbons, 1996; Lange, 2007; Arcidiacono et al. 2010; Kahn and Lange, 2014). The complex trade off that firms face between exploration and extraction has long been recognized as a key element of organizational learning (March, 1991) and a growing body of research combines insights from bandit problems[5] with statistical discrimination in contexts other than group learning (Che et al., 2019; Bardhi et al., 2020; Bergman et al., 2020; Fershtman and Pavan, 2020; Komiyama and Noda, 2020).

Like taste-based discrimination, my model generates differences between average performance and average pay of a group. In fact, it generates steady state predictions analogous to Becker (1957), with endogenous beliefs replacing preferences. Apparent taste-based discrimination can result from "incorrect" statistical discrimination and the model provides a new way to understand prejudice as the result of life experiences shaping beliefs in distortionary ways. Biased beliefs in my model differ starkly from a preference, highlighting that insights of prejudice-based models for labor market discrimination can be generated from uncertainty, without reliance on a utility function or biased updating.[6]

[4]Recent models in financial markets also consider agents who neglect the informational content of prices (Eyster et al., 2019), building on extensive evidence from voting, trading, investing, and auctions.

[5]See Bergemann and Välimäki (2008) for a review of bandit problems in economics.

[6]Individuals appear quick to form beliefs about groups and act on these in a way that shapes

Endogenous learning about groups differs from existing work on biased beliefs. Employers in my model are not fundamentally biased and attempt to maximize profits, but they conduct inference on a biased sample of observations about worker group productivity. The mechanism complements previous work on biased beliefs creating discrimination from true group differentials (Bordalo et al., 2016), biased updating (Sarsons, 2019), differences in the evaluation and supervision of workers (Bartoš et al., 2016; Glover et al., 2017), or implicitly (Bertrand et al., 2005). The model highlights that initial uncertainty about the relative productivity of worker groups leads to a learning problem that generates outcome differentials from biased beliefs. It provides a rationale for how even employers with no fundamental prejudice or bias, who are willing to give workers from any group a fair chance on profit maximization grounds, may endogenously develop persistent negative biases about the productivity of some groups.

This paper highlights that prejudice and statistical discrimination are not necessarily distinct or mutually exclusive, with implications for studying the source of discrimination (Bohren et al., 2019a; 2019b). The model generates different policy implications than previous models. For example, it provides a new lens to analyze affirmative action, which can induce employer learning by inducing minority hiring and improve outcomes as reported in Miller (2017). Consistent with evidence reviewed in Lang and Kahn-Lang Spitzer (2020) and in contrast with classical theories of discrimination, my model predicts that providing information on groups that is credible at the individual level can mitigate discrimination, as can encouraging intergroup interactions (Pettigrew and Tropp, 2006; Paluck et al., 2019).

future views, consistent with the notion of prejudice from psychology (Bertrand and Duflo, 2017). My model shows 1) how biases can micro-found the reduced-form notion of prejudice in economics and 2) how biases affect decision-making in statistical discrimination models. Beliefs are also particularly compatible with context-specific discrimination, such as variation across skill and education levels (Lang and Lehmann, 2012).

2.2 Labor Market Model

2.2.1 Employer Information and Beliefs

Consider a large number of employers hiring workers from two observably different groups A and B (e.g. race or gender). The key feature is that, through hiring, employers learn about the productivity of worker groups, which may differ across groups for example due to historical or social factors. Assume that employers know the productivity distribution of group A, but are initially uncertain about that of group B.[7] Information asymmetries across worker groups are a common feature in the literature, with the distinction that I focus on the dynamic implications of an initial asymmetry for hiring and learning (Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009; Lang and Manove, 2011).[8] Employers can learn about group B's productivity by hiring group B workers, but their objective is to maximize expected profits, leading to a potential trade off. Previous hiring experiences determine beliefs about the group's productivity and the value of additional learning.

Each worker, from either group, has productivity drawn from $X \sim N(\mu, 1/\tau)$.[9] For simplicity, assume that employers know the variance $1/\tau$ and that it is equal across groups. Employers know that group A's mean productivity is μ and have common priors about the mean productivity of group B, $\mu_B \sim N(\mu_0, 1/\tau_0)$.[10] I focus on the case where $\mu_0 = \mu$, such that employers have unbiased priors. Each employer hires one worker per period, uses their hiring experiences with group B to update their beliefs, and the match dissolves after each period.[11]

[7] The key feature is that initial information about group B's productivity is noisier, but assuming complete information on group A simplifies the analysis and exposition.

[8] Information asymmetry could arise in a majority versus minority setting where market participants naturally observe more information about the majority group over time. It could also arise if employers, for example from group A, have better information about workers of their own group due to previous experiences and interactions inside and outside the labor market.

[9] Appendix A extends the results to more general productivity distributions.

[10] Employers have misspecified beliefs, in the sense that groups are equally productive and the true mean productivity of group B μ is a fixed constant, but employers treat it as a random variable due to uncertainty.

[11] The implications of firm size are discussed in Appendix A. One-period contracts focus attention on group learning by studying employers repeatedly choosing between groups. Multi-period contracts may slow down learning if employers retain good workers, but they do not change relative incentives to hire and learn about group B, determined by μ_B.

I make three simplifications relating to hiring and learning. First, employers observe no individual signal of productivity prior to hiring; they rely solely on group membership to predict the productivity of a worker. Second, worker signals of productivity are private and only available through an employer's own hiring. Third, there is no human capital investment or signaling by workers. Each worker is endowed with a fixed productivity and inelastically provides a unit of labor each period. The implications of each simplification are discussed in Section 2.2.8 and Appendix A.

Workers hired from group B determine the information set of employer j, S_{jt}, composed of one private signal drawn from X for each hire. The cumulative number of signals employer j has observed by time t is $K_{jt} = \sum_{n=1}^{t} \mathbb{1}(L_{Bnj} = 1)$, where L_{Bnj} is an indicator variable for whether a group B worker was hired in period n. Employers form posterior beliefs about the mean group B productivity according to the Normal updating formula

$$\mu_B | S_{jt} \sim N\left(\frac{\tau_0\mu_0 + \tau\sum_{i=1}^{K_{jt}} x_i}{\tau_0 + \tau K_{jt}}, \frac{1}{\tau_0 + \tau K_{jt}}\right). \tag{2.1}$$

Letting $E[\mu_B|S_{jt}] = \frac{\tau_0\mu_0+\tau\sum_{i=1}^{K_{jt}} x_i}{\tau_0+\tau K_{jt}}$ and $\mathrm{Var}(\mu_B|S_{jt}) = \frac{1}{\tau_0+\tau K_{jt}}$, employers form posterior beliefs about group B productivity $X_B \sim N(E[\mu_B|S_{jt}], \mathrm{Var}(\mu_B|S_{jt})+1/\tau)$.[12]

2.2.2 Hiring Decision

Consider a frictionless labor market which clears each period. I first consider a model with infinitely-lived employers learning about one cohort of workers, abstracting from product-market competition through dynamic entry and exit of firms. Employers are risk neutral, wage-takers, and maximize the present value of lifetime profits. They consider the value of learning about the productivity of group B, leading to a dynamic optimization problem. An individual employer's posterior beliefs are characterized by $\psi_{S_{jt}} = \{E[\mu_B|S_{jt}], \mathrm{Var}(\mu_B|S_{jt})\}$ and Ψ_t is a list of posterior

[12]While the true variance in productivity $1/\tau$ is known, the posterior variance of X_B is larger since employers are uncertain about the mean, increasing expected variance. Formally, the variance is given by $\int \phi_{\mu_B|S_{jt}}(m) \int \phi(x|m)(x - E[\mu_B|S_{jt}])^2 dx dm = \mathrm{Var}(\mu_B|S_{jt}) + 1/\tau$.

beliefs across all employers. Group A's wage, w_A, is time-invariant and equal to their expected productivity μ. Group B's wage, $w_{Bt}(\Psi_t)$, is set competitively across employers through market clearing each period and evolves under the influence of Ψ_t. The current-period employer payoff from hiring a worker is simply equal to the productivity of their hire, x_i, with expected value μ for group A and $E[\mu_B|S_{jt}]$ for group B. Conditional on beliefs and wages at time t, employer j hires from group A or B to maximize their expected profits

$$
\begin{aligned}
V(\psi_{S_{jt}}, w_{Bt}(\Psi_t)) = \text{Max}\{&\mu - w_A + \beta E_t[V(\psi_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))], \\
&E_t[\mu_B|S_{jt}] - w_{Bt}(\Psi_t) + \beta E_t[V(\psi'_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))]\}
\end{aligned} \tag{2.2}
$$

where β is a discount factor. The continuation value $V(\cdot)$ includes updated beliefs $\psi'_{S_{jt+1}}$ when a group B worker is hired and $\psi_{S_{jt+1}} = \psi_{S_{jt}}$ otherwise. $E_t[V(\psi'_{S_{jt+1}}, \cdot)] \geq E_t[V(\psi_{S_{jt+1}}, \cdot)]$ since hiring from group B yields information which cannot decrease expected profits.

Endogenizing group B's wage is key because it is an outcome of interest, but also because intuition suggests that it should act as a counterbalancing force to biased beliefs. If the group's wage falls as a result of employers developing negatively-biased beliefs, then group B becomes relatively cheaper, which should in turn induce employers to hire them and learn, correcting biases. I study hiring and learning decisions which account for these endogenous wage adjustments.

Optimal hiring in the current period is determined by contrasting expected profits hiring from group B versus A. The difference is positive whenever

$$
\begin{aligned}
&\beta E_t[V(\psi'_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1})) - V(\psi_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))] > \\
&\mu - E_t[\mu_B|S_{jt}] - (w_A - w_{Bt}(\Psi_t)).
\end{aligned} \tag{2.3}
$$

Equation (2.3) compares the expected learning value from a group B hire on the left with expected foregone profit on the right. The perceived value of learning depends on the likelihood that it will lead to changes in hiring and higher expected profits.

It is maximized at $\mu_B = \mu$ since information is likeliest to affect subsequent hiring and decreases as μ_B becomes biased away from μ. In the case of negative bias, group B becomes less attractive from both a learning and a production standpoint. Thus, when prior experience suggests that group B is less productive, there is a trade off between expected learning benefits and expected foregone profits from hiring less productive workers. This trade off can be represented by a one-armed bandit problem, in which employers repeatedly choose between a "safe" arm (Group A) which yields a payoff from a known distribution and a "risky" arm (Group B) with an unknown payoff distribution. Obtaining comparatively low payoffs from the risky arm eventually leads the employer to stop experimenting and choose the safe arm, with the important distinction that wages and therefore payoffs are endogenous in my model.

One consideration is how employers learn about the productivity of group B from the evolution of their wage. In the baseline model, I rule this out by assuming static wage expectations: employers expect the wage next period to be equal to the current one, $E[w_{Bt+1}|S_{jt}] = w_{Bt}$, keeping the model tractable since employers do not form beliefs about the beliefs of other employers. The wage in theory does carry information relevant to the learning problem faced by employers. Yet, in practice, this assumption appears particularly mild given the complexity of the problem faced by employers. Market clearing wages summarize many private decentralized decisions that depend on factors unobserved by any given employer. Even if employers observe some relevant wage information and can invert the pricing and belief-updating processes, relative wages in practice are a function of many factors (changing skill and education, industry and occupation mixes, demographics, etc.), such that precisely isolating the impact of changing subjective employer beliefs appears implausible. Economists themselves have had long-standing unresolved debates about characterizing and decomposing wage gaps into components related to discrimination (Lang and Lehmann, 2012). Beyond recent work on financial markets which assumes that agents ignore the information value of prices supported by extensive evidence (Eyster et al., 2019), recent developments in modeling firm behavior surveyed in Aguirregabiria and Jeon (2019) focus on how uncertainty and learning in complex competitive environments can lead firms to have biased beliefs, for example about demand, costs, or the behavior of other firms.

Overall, taking the current wage as a prediction for the wage next period seems like a reasonable approximation in the context of the model, especially since it is correct in the long run. Still, deviating from this assumption by allowing employers to noisily learn about group B's productivity from their wage does not affect the qualitative predictions of the full model with dynamic market entry and exit of employers, as discussed in Section 2.2.8.

2.2.3 Hiring Cutoff and the Group B Wage

Define λ_{jt} as the relative willingness to pay (WTP) of employer j for a group B worker

$$\lambda_{jt} = \beta E_t[V(\psi'_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1})) - V(\psi_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))] - (\mu - E_t[\mu_B|S_{jt}]).$$

The trade-off between learning and foregone profit, ignoring wage considerations, is captured by λ_{jt} . It can be positive even if $E[\mu_B|S_{jt}]$ falls below μ, highlighting that employers may hire from group B even if they believe them to be less cost-effective to avoid future losses from incorrect beliefs.

Each period, labor market clearing implies that, at current wages, the fraction of employers who prefer to hire from group B is equal to the fraction of workers from the group. The group B wage each period is thus determined by the marginal employer m: the employer with the lowest λ_{jt} who must hire from the group to clear the market. Specifically, the wage is set such that the marginal employer is indifferent between hiring from either group, $\lambda_{mt} = w_{Bt}(\Psi_t) - w_A$, characterizing the optimal hiring strategy of employers stated in Proposition 1.

Proposition II.1 (Optimal Hiring)
The optimal hiring strategy of employers follows a cutoff rule where employer j hires from group B at time t if and only if $\lambda_{jt} \geq \lambda_t^c$. Moreover, $\lambda_t^c = w_{Bt}(\Psi_t) - w_A$.

Proof: See Appendix A.

Proposition 1 characterizes the cutoff below which it is optimal for employers to avoid hiring from group B at a given market wage, preserving their beliefs about the group's productivity. Since the wage gap is determined by $\lambda_t^c = \lambda_{mt}$, the optimal hiring decision of other employers immediately follows: those with λ_{jt} above the marginal hire from group B and others from group A, clearing the market. Market clearing thus implies the following condition

$$\nu_{\Psi_t}(\{\psi_{S_{jt}} : \lambda_{jt} \geq \lambda_t^c(w_{Bt}(\Psi_t))\}) = F_B \text{ and } \nu_{\Psi_t}(\{\psi_{S_{jt}} : \lambda_{jt} < \lambda_t^c(w_{Bt}(\Psi_t))\}) = F_A \tag{2.4}$$

where ν_{Ψ_t} is a measure over Ψ_t, F_g is the fraction of workers from group g, and each worker-employer pair has no incentive to deviate.

2.2.4 Equilibrium

An equilibrium is a stochastic process over beliefs and a mapping from beliefs to wages. Given a continuum of agents on each side of the market, this corresponds to a deterministic Markov process with corresponding transition functions characterized by Definition 1.

Definition II.1 *An equilibrium is a Markov process with a distribution over beliefs* Ψ_t *evolving according to a transition function* $T : \triangle(\mathbb{R} \times \mathbb{R}_+) \to \triangle(\mathbb{R} \times \mathbb{R}_+)$*, a wage function* w_{Bt}*:*

$\triangle(\mathbb{R} \times \mathbb{R}_+) \to \triangle\, \mathbb{R}$ *and an initial state* $\Psi_0 \in \triangle(\mathbb{R} \times \mathbb{R}_+)$ *such that every period:*

1. Employers make expected profit maximizing hiring decisions following equation (2.2)

and Proposition 1 for all $(\psi_{S_{jt}}, w_{Bt}(\Psi_t))$.

2. The labor market clears according to condition (2.4).

3. Employers update their beliefs:

a) Those with beliefs $\psi_{S_{jt}}$ *such that* $\lambda_{jt} < \lambda_t^c(w_{Bt}(\Psi_t))$ *hold posterior beliefs*

$\psi_{S_{jt+1}} = \psi_{S_{jt}}$.

b) Those with beliefs $\psi_{S_{jt}}$ such that $\lambda_{jt} \geq \lambda_t^c(w_{Bt}(\Psi_t))$ hold posterior beliefs $\psi'_{S_{jt+1}}$ derived according to equation (2.1).

The first condition states that employers maximize their expected profits according to their Bellman equation and the optimal hiring rule. The second condition states that the fraction of employers with beliefs such that they want to hire from group B given current wages (λ_{jt} above the marginal) is equal to the fraction of workers from group B. The third condition states that employers below the hiring cutoff for group B do not update their beliefs, while those above the hiring cutoff update their beliefs based on the productivity of their hire according to Bayes' rule.

2.2.5 Biased Beliefs and Discrimination

As a result of the optimal hiring rule and equation (2.1), it is straightforward to characterize the asymptotic distribution of posterior beliefs described in Proposition 2.

Proposition II.2 (Asymptotic Beliefs and Persistent Negative Biases) *As $t \to \infty$, beliefs of employers who remain above the hiring cutoff converge in distribution to μ. Others hold a range of beliefs such that $E[\mu_B|S_{jt}] < \mu$. The limiting fraction of employers with $E[\mu_B|S_{jt}] < \mu$ equals the fraction of group A workers.*

Proof: See Appendix A.

By standard Bayesian reasoning, posterior beliefs converge to the truth as the number of signals goes to infinity. On the other hand, employers below the cutoff (which implies $E[\mu_B|S_{jt}] < \mu$ in the long run given a strictly positive value of learning) do not hire from group B, preserving negative biases. In the long run, since unbiased employers hire from group B and biased employers hire from group A, the fraction of biased employers is equal to the fraction of group A workers.[13] Proposi-

[13] The Becker (1957) taste-based model requires that the fraction of prejudiced employers be at least as large as the fraction of group A workers to generate a wage gap. Both models thus require a majority of biased or prejudiced employers to generate a wage gap if group A is larger than group B.

tion 2 highlights that optimal hiring and learning lead a subset of employers to hold negatively-biased beliefs, even asymptotically.

Endogenous employer learning about worker group productivity generates a plausible distribution of beliefs for discrimination to arise. First, beliefs about group B's productivity exhibit sustained heterogeneity across employers. Second, differential learning across employers results in beliefs being negatively-skewed. The endogenous learning mechanism generates these features without relying on group differentials,[14] prejudice, or biased priors, providing a novel way to understand persistent, heterogeneous, negatively-biased beliefs about a worker group's productivity.

The next consideration is whether these biased beliefs generate discrimination in the form of a wage gap. Proposition 3 characterizes the evolution of group B's wage.

Proposition II.3 (Wage Gap and Persistent Discrimination)
$w_{Bt}(\Psi_t)$ is strictly decreasing in t and converges to a constant $c < w_A$.

Proof: See Appendix A.

The distribution of beliefs becomes negatively-skewed with time, because only negative bias can be stable. With hiring experience, supramarginal values of λ_{jt} become concentrated around 0 as $E[\mu_B|S_{jt}]$ becomes concentrated around μ. By definition, λ_{mt} lies below supramarginal values of λ_{jt} and thus eventually falls below 0, leading $w_{Bt}(\Psi_t)$ to fall below w_A. By market clearing, the wage cannot increase or remain constant with time. Given a continuum of employers, some employers just above the hiring cutoff are expected to have relatively negative hiring experiences with group B in any given period, such that their λ_{jt} fall below that period's cutoff. Then, the fraction of employers who prefer to hire from group B at the current wage becomes lower than the fraction of group B workers. The wage must thus decrease to induce

The fraction of employers with biased beliefs in my model is endogenously determined to be exactly equal to that of group A by market clearing, rather than being assumed. Widespread biased beliefs may be more plausible than widespread animus, and Lang and Lehmann (2012) discusses evidence that a large share of employers hold negative perceptions in the context of race. Moreover, Black (1995) shows that wage gaps may be sustained under milder conditions in a search framework, as briefly discussed in Appendix A.

[14]This distinction has important implications even when it is unlikely that two groups have equal productivity in practice, since it predicts that closing productivity gaps would not necessarily eliminate discrimination.

employers to hire from the group and clear the market. Lastly, since beliefs are fixed asymptotically, there is virtually no updating and no change in the wage, so it converges to a constant.

Since both groups are equally productive, the wage gap implies that group B is paid below their expected productivity. While the predicted wage gap depends on relative group productivity, the prediction that group B is paid below their expected productivity does not. The model thus predicts that persistent negatively biased employer beliefs about group B's productivity arise endogenously through hiring interactions and generate persistent discrimination against the group.

2.2.6 Entry, Exit and Competition

A common view is that market competition should drive out biased beliefs and therefore resulting discrimination, at least in the long run. To investigate this, I augment the model with dynamic employer entry and exit from the market. The fundamental intuition regarding differential learning across employers and therefore biased beliefs remains, but exit provides a straightforward reduced-form way to introduce competition through differential exit rates based on beliefs.

Employers exit the market and are replaced with new employers at an expected aggregate rate δ each period. The exit rate influences the expected duration in the market, learning incentives, and available time for employers to potentially correct their biases. It can also directly affect the belief distribution by introducing new employers who hold different beliefs on average. I assume that employers enter with unbiased priors, although Appendix A shows that discrimination can be amplified when priors are influenced by experienced employers.[15]

Hiring and wage determination follow the same process as before. Profit maximiza-

[15] Prior variance may decrease if employers learn from previous cohorts of employers. This is unlikely to eliminate the initial information asymmetry since it would require employers to completely ignore their experiences, going against evidence that decision-makers put too much weight on their own information (Moore et al., 2015), and because the learning problem in practice is constantly changing across cohorts of workers, such that employers must rely on their experiences to assess group productivity in their own context. For example, the relative education and experience of women and minority workers compared to that of white men was not the same in 1990 as it is today, and employment contexts have changed substantially.

tion is given by

$$
\begin{aligned}
V(\psi_{S_{jt}}, w_{Bt}(\Psi_t)) = \text{Max}\{&\mu - w_A + (1-\delta)\beta E_t[V(\psi_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))], \\
&E_t[\mu_B|S_{jt}] - w_{Bt}(\Psi_t) + (1-\delta)\beta E_t[V(\psi^{'}_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))]\}.
\end{aligned}
$$

The exit rate of an employer should depend on profits and therefore hiring decisions determined by $E_t[\mu_B|S_{jt}]$. Since $E_t[V(\psi^{'}_{S_{jt+1}}, \cdot)] \geq E_t[V(\psi_{S_{jt+1}}, \cdot)]$, employers who hire from group B earn higher expected profits of at least $w_A - w_{Bt}$ each period. Given a lower wage and equal productivity for group B, these employers are more profitable and accordingly should have a lower market exit rate, $\delta_B < \delta_A$ with $\delta = \delta_B F_B + \delta_A F_A$. If the only determinant of market exit is beliefs about the productivity of group B ($\delta_B = 0$), a differential exit rate eliminates discrimination at least in the limit.[16] Yet, firm survival in a market depends on many factors, such that firms who hire from group B also exit the market and biased beliefs may often not be pivotal (Audretsch, 1991; Schary, 1991; Black, 1995; Hellerstein et al., 2002).

The key point is that biased beliefs are not a primitive of the model, but arise endogenously. Therefore, as some employers held unbiased priors but developed biased beliefs through hiring, so may new employers. In the aggregate, biased beliefs and the wage gap are not necessarily eliminated by competition. Depending on parameters, a wage gap can be sustained asymptotically even if employers who hire from group A are driven out at a higher rate, as summarized in Remark 1.[17]

Remark 1 (Persistent Discrimination with Market Competition)
For some values of δ_A and δ_B with $\delta_A > \delta_B$, there exists a period $\bar{t}$ in which $w_{Bt}(\Psi_t)$ falls below w_A, remains below for all $t > \bar{t}$, and converges to a constant $c < w_A$.

Remark 1 is illustrated through simulation in the next subsection. The main diffe-

[16] Beliefs of employers above the hiring cutoff for group B converge to the group's true productivity, so an arbitrarily small mass of new entrants with $\lambda_{jt} \geq 0$ guarantees that w_{Bt} is not below w_A.

[17] In taste-based models, firm growth is important since prejudiced firms may remain in the market earning lower profits to indulge in their taste for discrimination. Then, discrimination is mitigated because unprejudiced firms grow more quickly. In my model, firms are not willing to accept a lower return for their mistaken beliefs, so growth is not conceptually necessary for discrimination to be competed away.

rence with Proposition 3 is that the existence of a wage gap depends on parameters. At one extreme, for exit rates near zero, the existence of a wage gap directly follows from Proposition 3. At the other extreme, for very high exit rates, it is possible to introduce enough new employers with unbiased priors to hire all of group B each period, eliminating the wage gap. At the intensive margin, higher competition reduces the magnitude of the wage gap as shown in Appendix A and consistent with empirical evidence (Ashenfelter and Hannan, 1986; Black and Strahan, 2001). At the extensive margin, competition may not eliminate discrimination arising from endogenous biased beliefs.

2.2.7 Simulations

To illustrate the model's dynamics, a set of simulations was computed over 1,000 periods with 10,000 employers and 10,000 workers, 25% of which are from group B. Simulation details are outlined in Appendix A, along with additional results. Because the simulated market is finite, the evolution of beliefs and wages is stochastic rather than deterministic. Emphasis should be put on the model dynamics characterized by Propositions 1-3 and Remark 1, which do no substantively vary with parameter choice, rather than specific values of the wage gap.[18]

Panel A of Figure 2.1 shows the evolution of beliefs for key moments of the distribution, without entry and exit. The 25% of employers with the highest valuation for group B each period hire them and learn, so their beliefs converge towards the group's true mean productivity normalized at 0, while those of other employers are negatively biased and do not evolve. Panel B shows that the group B wage initially lies above the marginal employer's beliefs due to the value of learning, but eventually falls and remains below zero (also normalized as the group A wage) as beliefs fall below μ and the value of learning falls. With a finite market, there is a separation in the WTP of employers above and below the cutoff, seen in Panel A between the 75th and 76th percentiles. The market clearing wage can lie anywhere between these

[18]Similarly, the initial state in which employers enter the market exhibits theoretically intuitive features, but is of limited practical interest. Given all employers entering simultaneously with unbiased priors, the initial group B wage may be higher than that of group A because of market clearing, but this depends on prior beliefs, relative uncertainty and productivity across groups, and potential ambiguity aversion.

two percentiles, while the latter determines the wage with a continuum of employers as characterized in Proposition 3. If match surplus is allocated to employers, the wage is also set by the 76th percentile with a finite number of employers, as shown in Panel B.

Similarities and differences between the simulated wage path and empirical wage trends, namely whether the wage is increasing, stable, or decreasing over time, naturally do not provide a test of the model's key implications. Empirical trends depend on many sources of wage differentials outside of the model, while simulated trends depend on assumptions on priors and relative productivity, among others. For example, Appendix A shows that negatively-biased priors can generate a group B wage which starts and remains below that of group A, but increases over time. An analogous argument can explain the seemingly odd model prediction that employers begin by hiring group B most often and gradually decrease their hiring of the group, rather than potentially the other way around.

Figure 2.2 presents simulations with market entry and exit, a 2% aggregate exit rate each period, and a 25% higher exit rate for employers below the hiring cutoff. The set of employers in the market is expected to be jointly replaced 3 to 4 times over the period, so the pattern is simply repeated beyond. One notable difference is that, since all employers exit the market in finite time, some employers above the hiring cutoff may always have negatively-biased beliefs. There is thus a sense in which entry and exit can actually help sustain a wage gap by preventing belief convergence.

2.2.8 Outside Learning

If employers observe information about group B's productivity outside of their own hiring, such as the hiring decisions or outcomes of a competitor, the performance of group B in other settings, or the evolution of wages, they may learn without hiring. Such outside learning can mitigate or exacerbate bias, but has limited impact on the model's key predictions.

Consider a benchmark case in which employers get one outside signal about group B productivity per period irrespective of hiring. Outside signals are distributed $O \sim N(\mu, 1/\tau_o)$. Posterior beliefs are given by

$$\mu_B|S_{jt} \sim N\left(\frac{\tau_0\mu_0 + \tau\sum_{i=1}^{K_{jt}} x_i + \tau_o\sum_{m=1}^{t} o_m}{\tau_0 + \tau K_{jt} + \tau_o t}, \frac{1}{\tau_0 + \tau K_{jt} + \tau_o t}\right). \tag{2.5}$$

There are several reasons to expect that employers put more weight on their own signals, $\tau_o < \tau$. Even at similar firms, there is some degree of mismatch between employment contexts. It may be difficult for an employer to learn about the productivity of group B from observing others when hiring and performance depend on many factors beyond employer beliefs. The employer learning literature indicates that learning is asymmetric; employers have better information on their hires than other employers do (Waldman, 1984; Kahn, 2013; Ge et al., 2020). Employers also have a tendency to over-weight their own information (Moore et al., 2015). Chapter IV reports in the context of a large national retailer that a manager's hiring of black workers is influenced by their own previous hiring experiences with the group, but not those of other managers even within the same store. Similarly, recovering a signal about group B's productivity from wage information that is relevant to a given employer's specific hiring context is likely to be particularly difficult. In short, it's not clear what form outside information would need to take to be credible at the individual level to an employer who has already formed beliefs based on their own experience.

If employers put more weight on their own signals, then those who hire group B still learn faster, especially if they also observe outside signals. The belief distribution remains negatively-skewed in any finite period, and the bias-generating mechanism at the least slows down learning. Slowing down learning itself has non-negligible implications. Statistical discrimination generally predicts that the market immediately learns equilibrium worker group productivity. One criticism is that learning is "too fast" for these models to be important in the long run (Lang and Lehmann, 2012). My model explains why learning about some groups may be particularly slow and create discrimination along the equilibrium path, reducing the lifetime income of these groups.

In the long run, if beliefs converge over time, then the wage gap is eliminated. If beliefs do not fully converge, for example because there is market entry and exit or the learning problem evolves over time, then the wage gap can remain. In practice, these two conditions are clearly satisfied. Firms, employers, and recruiters regularly enter and exit the market, and the relative productivity of worker groups has been evolving with changes in demographics and education, among other factors. Remark 2 summarizes this result for the case of market entry and exit, which again follows from Proposition 3.

Remark 2 (Persistent Discrimination with Outside Learning)
For some values of τ_o, τ, δ_A, and δ_B with $\tau > \tau_o$ and $\delta_A > \delta_B$, there exists a period $\tilde{t}$ in which $w_{Bt}(\Psi_t)$ falls below w_A, remains below for all $t > \tilde{t}$, and converges to a constant $c < w_A$.

Moreover, outside information also poses some challenges. For instance, making hiring outcomes public within employer networks does not conceptually solve the issue that employers learn too little, because it lowers incentives for employers to hire group B and learn from their own signals, leading to free-riding (Keller et al., 2005; Hoelzemann and Klein, 2018). Equation (2.5) also assumes that outside signals are unbiased, unambiguous, and unrelated to existing bias. Otherwise, outside signals could preserve or exacerbate biased beliefs (DeGroot, 1974; Gentzkow and Shapiro, 2006; Baliga et al., 2013; Enke and Zimmermann, 2017; Fryer et al., 2018). In any case, outside learning suggests two implications. First, discrimination may differ across settings based on the observability of competitors, workers, wages, and output. Second, there is potential scope for the design and provision of information.

2.3 Relationship with Other Theories

The model generates steady state predictions analogous to those from Becker (1957), with preferences replaced by endogenous beliefs:

- An employer hires group A if the wage gap is smaller than λ_{jt} and group B otherwise.
- If enough employers have (approximately) correct beliefs to hire all of group B, there is effective segregation without wage gap.
- If enough employers have biased beliefs, there is a wage gap determined by the

marginal employer.

The model has intuitive similarities with taste-based discrimination, namely a difference between average productivity and average pay of a group, but without deviating from a statistical discrimination framework. This is a key point given that taste-based discrimination has often been criticized for the arbitrariness of including preferences in a utility function. The important insights of prejudice-based models for labor market discrimination do not in fact rely on preferences, but can be understood as arising from uncertainty. Biased beliefs capture context-dependent aspects such as gender-based discrimination and differentials by skill and education, which are less compatible with the notion of an aversion to contact. Widespread biased beliefs may also be more plausible than widespread overt animus, which evidence suggests has been steadily decreasing over past decades, unlike outcome differentials (Lang and Lehmann, 2012). This does not imply that preferences and biased beliefs are necessarily substitutes, because they differ fundamentally in how discrimination arises, evolves, and can be mitigated.[19]

The model complements the statistical discrimination literature by relaxing the assumption that employers have correct equilibrium beliefs about group productivity and instead modeling learning. In many contexts, the assumption that employers know the productivity of worker groups or instantly learn it in equilibrium seems implausible, yet little work considers how relaxing the assumption can have important implications.[20] Discrimination caused by biased beliefs can arise without grounds for classical statistical discrimination. It does not arise from employers using *objective* information about groups, but their potentially flawed beliefs. It is not a self-fulfilling prophecy nor the result of coordination failures between firms and workers. The discriminated-against group cannot be seen as having "played a hand" in justifying discrimination against them, and discrimination can be sustained without prior bias or homogeneous beliefs.[21]

[19] If biased beliefs are reinforced through behavioral primitives in the utility function, they could be essentially indistinguishable from a taste. Individuals with a taste for discrimination may gather and interpret information in a way that validates and justifies their prejudice (Nickerson, 1998).

[20] Aigner and Cain (1977) state in their model that group means "are estimated without bias" by employers and that "as an explanation of discrimination against blacks, a theory of discrimination based on employers' mistakes is even harder to accept than the explanation based on employers' 'tastes for discrimination,' because the 'tastes' are at least presumed to provide a source of 'psychic gain' (utility) to the discriminator."

[21] The homogeneous prior assumption usually made in the literature can be important to generate

Another point concerns efficiency and equality. In statistical discrimination models, outcomes usually reflect true average productivity, so ending discrimination may not help group B on average. As a result, this type of discrimination is generally regarded as efficient. In my model, workers are paid below their expected productivity because of what are essentially employer mistakes. A social planner concerned with inequality or equality of opportunity could improve group B outcomes at no efficiency cost through increased employer learning.

2.4 Implications for Empirical Work and Policy

Discrimination from biased beliefs is consistent with a growing body of evidence. It has implications for identifying the source of discrimination, which has traditionally meant distinguishing between taste-based and statistical discrimination. Empirical tests often provide indirect evidence by comparing observed outcomes to those expected from true group differences, with the residual classified as taste. Such logic is conceptually inadequate, because the absence of observable productivity differentials does not imply a taste for discrimination given statistical discrimination with incorrect beliefs. Similarly, employers responding to information is consistent with statistical discrimination, but does not imply that employers hold correct beliefs on average or use information correctly. Bohren et al. (2019a) studies the empirical identification challenge posed by biased beliefs and stresses that they are rarely considered in the literature. My model provides a new bias-generating mechanism in the labor market and blurs the line between the two classical theories, highlighting that biased beliefs should not be ignored as a potential source of discrimination.

Identifying the source of discrimination is important partly because policy implications can differ. Competition can mitigate discrimination in the case of taste or biased beliefs, but may not eliminate it if information asymmetries remain. Closing productivity gaps may mitigate discrimination based on true group differentials, but those are not necessary for belief-based discrimination to persist. Diversity or implicit bias training could provide relevant information about groups, but if they target

long-run discrimination. Otherwise, some employers may be better at interpreting signals (Aigner and Cain, 1977) or have more accurate priors (Coate and Loury, 1993). Other employers would learn or exit the market, such that the need for the discriminated-against group to adjust is unclear.

cognitive biases and implicit stereotypes, will not address biased beliefs as in my model. Providing information on individual productivity may mitigate statistical discrimination by decreasing reliance on group membership, but information on productivity may help distinguish between the different theories. Information about groups should have little impact if employers have correct beliefs on average or if animus is driving discrimination, but may mitigate biased beliefs consistent with mounting evidence surveyed in Lang and Kahn-Lang Spitzer (2020). As mentioned previously, information about groups must be perceived as informative to an employer who has potentially already formed beliefs based on their own experience. Accordingly, policies which induce individual employers to learn more through their experiences may be particularly effective.

Indeed, central to the model is the idea that employers learn about groups through interaction and exposure. My model provides a new lens to study policies like internships, worker subsidies, and affirmative action which can push employers to hire more workers from group B and learn, consistent with improved minority outcomes as documented in Miller (2017). Critics of affirmative action often state that the worker best qualified for a position should be hired, independent of group membership. This argument hinges on the assumption that employers know ex-ante which worker is most qualified and therefore have correct beliefs about group productivity. My model suggests that this may not be the case and that affirmative action may in fact be necessary to move towards the point where the worker best qualified for a position is hired, independent of group membership. Relatedly, Pettigrew and Tropp (2006) and Paluck et al. (2019) conclude from their surveys that intergroup contact, particularly intense collaborative exposure and integration, typically reduces prejudice. These predictions follow directly from my framework of belief updating. One historical example is World War II, often discussed as a shock through which employers learned about the productivity of women and minority groups (Goldin, 1991).

Consistent with endogenous employer learning about groups, Leung (2017) documents on an online job board that previous hiring experiences of employers with workers from particular countries affect the subsequent likelihood of hiring workers from those countries. Chapter IV uses longitudinal employment records from a large US retailer and documents that the hiring history of managers creates heterogeneity in their hiring of worker groups. Managers increase (decrease) their relative hiring of black and white workers following positive (negative) experiences with these groups, with proportionally larger impacts for black workers consistent with stronger updating by managers. Further, early negative experiences with black workers persistently decrease relative hiring of the group over subsequent hiring cycles, unlike early positive experiences or early negative ones with white workers. These findings are particularly consistent with the idea that hiring experiences of employers lead them to update their beliefs about the performance of worker groups, affecting subsequent hiring patterns in a manner consistent with endogenous employer learning systematically decreasing relative hiring of minority workers. These papers suggest that studying how individual discriminatory responses evolve over time is a key avenue to test belief-based discrimination and distinguish it from other sources.

Regarding HR policy, firms with decentralized hiring in which individual managers hold discretionary power may especially have incentives to eliminate biased beliefs as studied in this paper. Hoffman et al. (2018) finds that managers who hire against job-testing technology recommendations tend to hire worse workers, consistent with mistakes or biases. Berson et al. (2019) suggests that discrimination among large firms appears lower at firms with centralized hiring. Bergman et al. (2020) study resume screening algorithms in a setting where firms balance selecting workers from previously successful groups with selecting from under-represented groups. They find that algorithms which value learning can improve both hiring performance and diversity, suggesting that policies which gather and share information within firms can help mitigate the impact of individual biased beliefs.

2.5 Conclusion

This paper presents a new statistical discrimination model in which persistent, heterogeneous employer biased beliefs about the productivity of worker groups arise and can create discrimination. Given initial uncertainty about the relative productivity of worker groups, employers systematically develop biased beliefs through endogenous learning decisions influenced by their previous hiring experiences with groups. These biased beliefs can create discrimination against worker groups whose productivity is initially more uncertain to employers, even with expected profit-maximizing employers in a competitive market with equally-productive worker groups, no prior bias or prejudice, and without endogenous worker investments.

Empirical evidence from an experimental labor market designed to test the individual-level implications of endogenous employer learning on hiring and beliefs supports the model's core mechanism. Positive experiences with a worker group lead to more hiring and learning, correcting positive biases. Negative experiences decrease hiring and learning, preserving negative biases and leading to a negatively-skewed employer belief distribution about the productivity of a worker group whose productivity is initially unknown.

The model generates steady state predictions analogous to Becker (1957), replacing preferences with endogenous biased beliefs and highlighting that some of what is usually classified as a taste may be understood as biased beliefs. It provides a new way to understand prejudice in the labor market as the result of selected interactions between groups distorting beliefs and behavior. It generates these novel implications while being set within a statistical discrimination framework in which learning about groups is modeled explicitly, complementing previous models in that literature. Biased beliefs in this paper arise from information frictions, with implications for understanding the relationship between theories of discrimination, empirically studying the source of discrimination, and policy.

The model focuses on profit-maximizing employers who are Bayesian over their own experiences, although existing work documents behavioral elements which could amplify discrimination based on biased beliefs and increase the connection with preferences. This interaction is a natural direction for future research and suggests that biased beliefs in this model may constitute a lower bound in many empirical settings.

Figure 2.1: Model Simulation without Entry and Exit

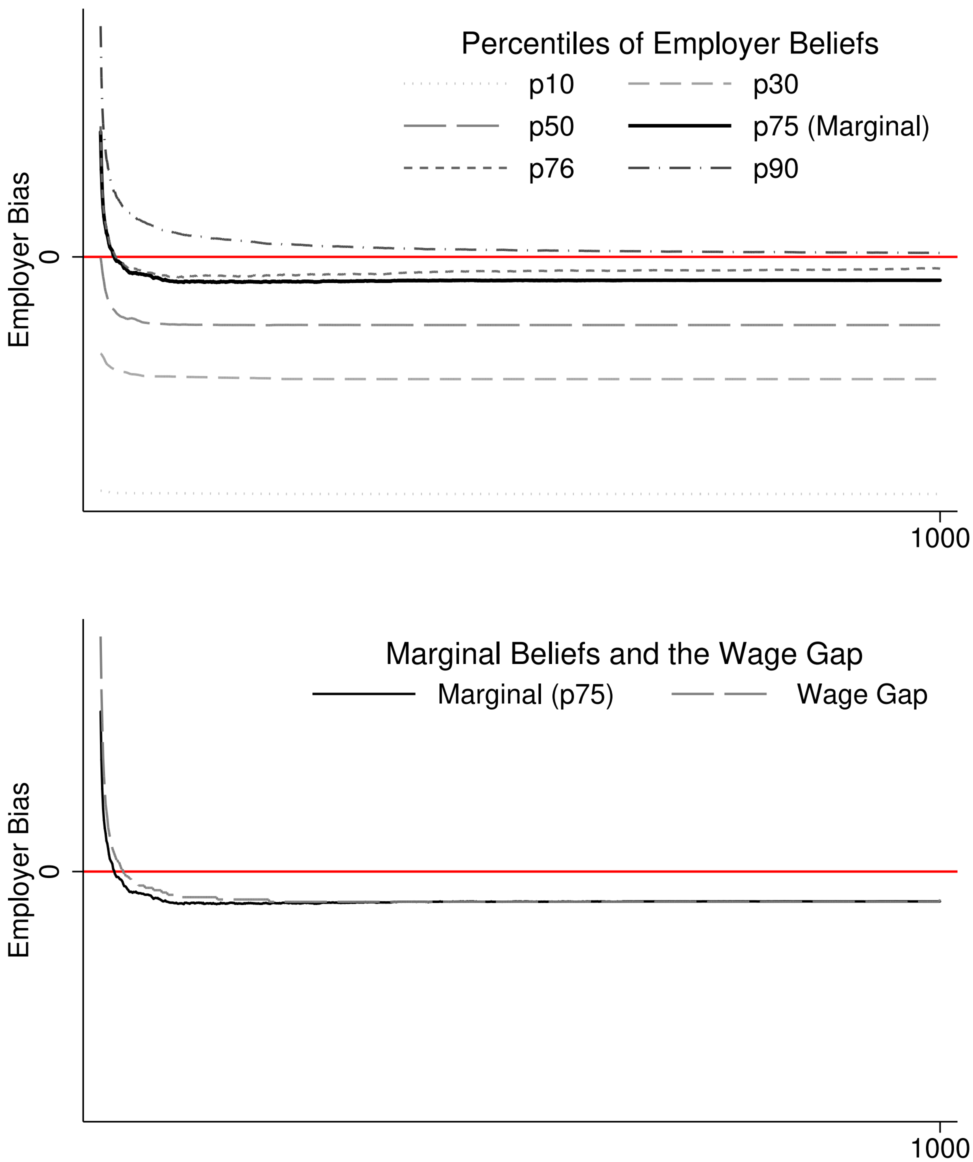

The fraction of group B workers is 0.25. Worker productivity is distributed $N(0, 2)$, prior beliefs are distributed $N(0, 1)$. w_A is normalized to 0 and β is set to 0.9.

Figure 2.2: Model Simulation with Market Entry and Exit, 25% Exit Differential

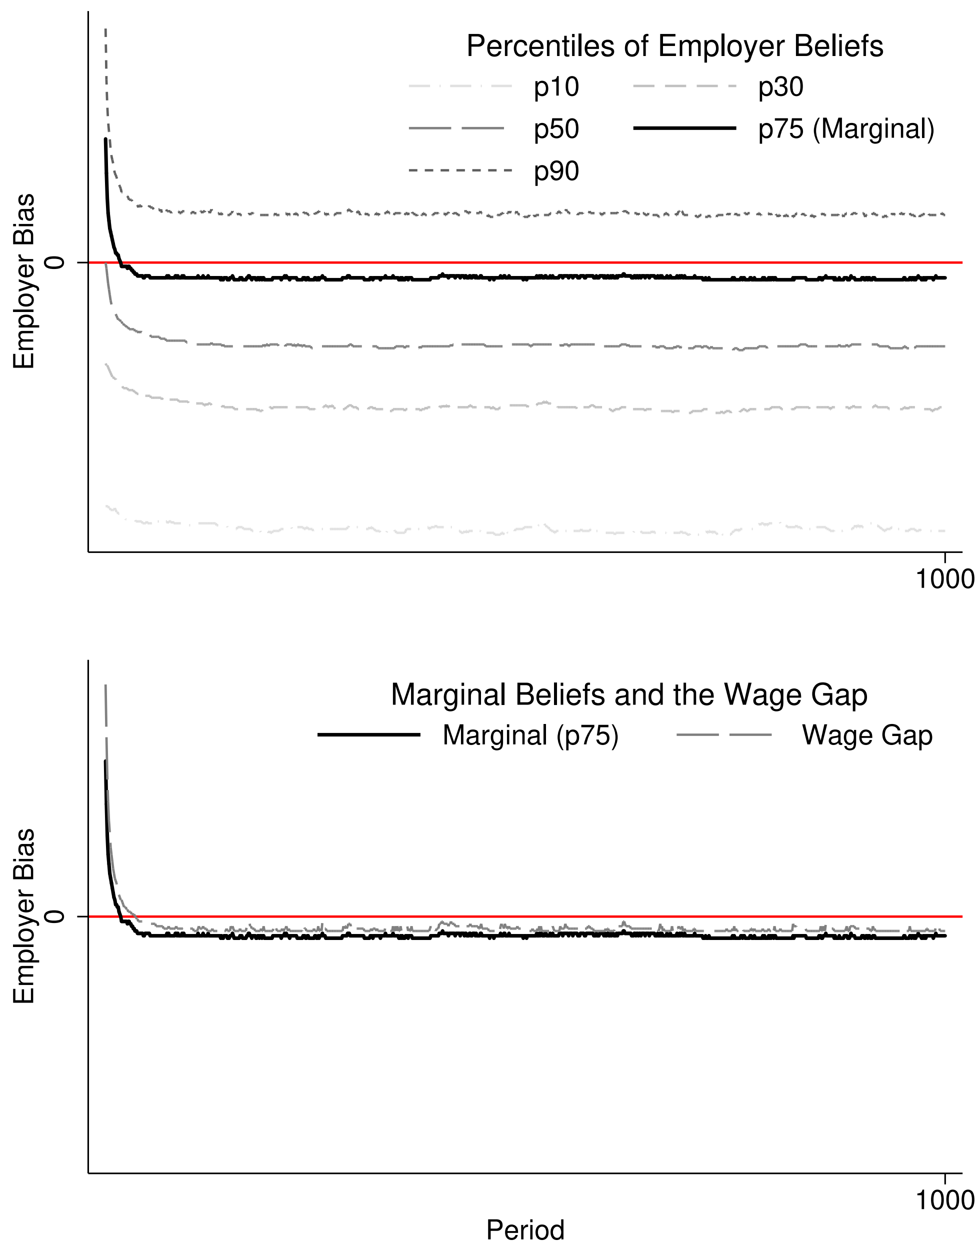

The aggregate exit rate corresponds to 2% each period, with a 25% higher exit rate for employers below the hiring cutoff for group B. New entrants have mean beliefs equal to 0 (unbiased). See Figure 2.1 for other parameter choices.

CHAPTER III

Experimental Evidence on Endogenous Stereotype Formation and Hiring Discrimination

3.1 Introduction

Evidence across the social sciences suggests the existence of pervasive negative perceptions against certain groups of workers in the labor market (Kirschenman and Neckerman, 1991; Holzer, 1996; Wilson, 1996; Biernat and Kobrynowicz, 1997; Moss and Tilly, 2001; Pager and Karafin, 2009; Pager et al., 2009). In economics, a growing literature studies the role of negative perceptions, as potentially biased or incorrect beliefs about groups, in generating discrimination (Fershtman and Gneezy, 2001; Bertrand et al., 2005; Reuben et al., 2014; Bordalo et al., 2016; Mobius et al., 2016; Arnold et al., 2018; Bohren et al., 2019a; 2019b; Bordalo et al. 2019; Sarsons, 2020). Yet, little work focuses on specific mechanisms through which biased beliefs can arise, persist, and generate discrimination in a labor market setting.[1]

This chapter uses an experimental labor market to test the theoretical mechanism proposed in Chapter 2 through which negatively-biased beliefs about the productivity of worker groups arise endogenously through hiring experiences of employers with workers. The intuition for the mechanism is as follows. If employers are uncertain about the productivity distribution of worker groups when they begin hiring

[1] For example, a standard assumption in the statistical discrimination literature is that employers learn about the productivity of individual workers, but not that of groups. They either have complete information about groups (Aigner and Cain, 1977) or are assumed to have correct equilibrium beliefs about group productivity (Coate and Loury, 1993).

workers, then, through their own hiring experiences, employers learn about both the productivity of their hires and that of their group. The nature of hiring experiences with a group then determine subsequent incentives to hire from the group again and learn more about their productivity, generating differential hiring and learning across employers. Positive experiences create positive biases which correct themselves since they lead employers to hire from the group again and learn about their productivity. Negative experiences lead employers to reduce or stop hiring from the group on expected productivity grounds, decreasing learning, preserving negative biases, and creating a negatively-skewed belief distribution about the group's productivity. An employer's hiring history with a group thus endogenously shapes their subsequent hiring and learning, disproportionately impacting workers from groups whose productivity is initially more uncertain to employers, such as minority groups (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and Várdy, 2009; Lang and Manove, 2011).

Crucially, the mechanism recasts biased beliefs as the systematic result of expected profit maximization and Bayesian updating of employers over their own experiences in a setting of initial uncertainty about the productivity distribution of worker groups. It does not rely on taste-based discrimination (Becker, 1957), productivity differences, self-fulfilling prophecies, or biased employer priors. The mechanism departs from those based on the distortion of true group differences (Bordalo et al., 2016), bias in the evaluation of workers (Bartoš et al., 2016), or bias in the belief updating itself (Sarsons, 2020), highlighting a distinct complementary source of bias. It reflects the notion that a substantial component of the information that an employer observes through their hiring is privately observed and that an employer's own experiences may play a particularly important role in shaping their beliefs and behavior (Waldman, 1984; Kahn, 2013; Moore et al., 2015; Ge et al., 2020; Guenzel and Malmendier, 2020). It also provides a rationale for how even employers with no fundamental bias, who are willing to give workers from any group a fair chance on profit maximization grounds, may endogenously develop persistent negative biases about the productivity of some groups. As a result, biased beliefs may be more widespread, pervasive, and resistant to market forces than typically understood.

I create a controlled environment to study how biased beliefs arise through endogenous employer learning about groups. I specifically abstract away from existing real life biases and discrimination by creating minimal worker groups where membership is randomly assigned to one of two color groups.[2] Workers perform the real-effort task of solving character puzzles on the computer, which corresponds to their productivity in the context of the experiment. Employers then repeatedly hire workers, choosing between one of the two groups each period and observing their hire's productivity. They are incentivized to hire the most productive workers available, requiring them to identify which group is more productive, if any. I study how negative biases about group productivity arise from an employer's hiring history. I give employers better initial information on the productivity of one group and focus on hiring and learning about the other group. I elicit employer beliefs about the group's productivity after each time they hire a worker from the group, allowing me to track the evolution of biased beliefs and their impact on subsequent hiring behavior.

The results show that negative hiring experiences with the group whose productivity is initially more uncertain, captured through the hiring of relatively low productivity workers, lead to persistent negatively-biased beliefs about the group's productivity, specifically by decreasing subsequent hiring of the group and therefore learning. In contrast, positive experiences increase subsequent hiring and learning, mitigating positive biases. Across employers, differential hiring and learning generate a persistent negatively-skewed distribution of beliefs about the group's productivity. In fact, compared to their initial belief distribution, employers on average have no more accurate beliefs about the group's productivity after the hiring experiment. Using additional experimental treatments, I also show that the specific hiring context matters for the formation of biased beliefs. In particular, evidence suggests that bias formation against a group is particularly strong when it is presented as a minority group and when employers are not primed to think about their beliefs until the end of the experiment. Lastly, I also provide evidence that employers update their beliefs by more than a Bayesian benchmark following their experiences, consistent with stereotype formation amplifying the effects of the endogenous learning mechanism in practice.

Experiments have frequently been used to study features of discrimination in econo-

[2] Using minimal groups based on arbitrary characteristics is particularly valuable to isolate the mechanism by abstracting from existing beliefs associated with groups like gender or ethnicity (Tajfel et al., 1971; Charness et al., 2007; Chen and Li, 2009; Sutter, 2009; Goette et al., 2012).

mics (Charness and Kuhn, 2011; Neumark, 2018). They are particularly well-suited to study belief-based discrimination because they provide a rare environment in which beliefs can be observed and mapped into behavior. Fershtman and Gneezy (2001), Mobius et al. (2006), Mobius et al. (2016) and Bohren et al. (2019a; 2019b) provide experimental evidence of discriminatory behavior driven by biased beliefs in the case of gender, beauty, and ethnicity. Dianat et al. (2018) conducts an experiment in which differences in human capital investment across worker groups lead to statistical discrimination which persists partly because employer beliefs about groups are slow to adjust. In contrast, rather than documenting existing biases against specific groups, I focus on how biases arise endogenously in a hiring setting with no grounds for discrimination based on productivity differentials to arise and explicitly distinguish biased beliefs from taste-based discrimination. Since initial experiences play a particularly important role in determining future hiring and beliefs, the mechanism relates intuitively to work on the lasting consequences of first impressions (Olivora and Todorov, 2010; Agnew et al., 2018; Oh et al., 2020).

3.2 Theoretical Framework

Chapter 2 presents a formal model of the mechanism and hypotheses presented in this section. Consider two groups of workers *Gray* (G) and *Orange* (O), both with i.i.d. productivity $X \sim N(\mu, 1/\tau)$. Employers are initially uncertain about potential group differences in productivity and thus perceive group membership as a relevant indicator of productivity.[3]

For expositional simplicity, assume that employers know the distribution of productivity for group O and the variance in productivity for both groups. This corresponds to a one-armed bandit problem: employers hire from group O (safe arm) with known productivity distribution or Group G (risky arm) with unknown mean productivity.[4] Employers have priors about the mean productivity of group G given by

[3] In practice, employer estimates of productivity also depend on individual signals like resumes, but these have little impact on the mechanism's qualitative implications, as discussed in chapter 2.

[4] See Bergemann and Välimäki (2008) for an overview of bandit problems. They have been implemented in experiments studying whether participants follow optimal strategies, showing that participants value learning but switch between arms too often and learn less than optimal (Meyer and Shi, 1995; Banks et al., 1997). One explanation for this last finding is ambiguity aversion

$\mu_G \sim N(\mu_0, 1/\tau_0)$ and learn about the group through hiring. Hires from group G determine the information set of employer j at time t, S_{tj}, composed of one private signal per hire. Employers form posterior beliefs about the mean productivity of group G

$$\mu_G|S_{tj} \sim N\left(\frac{\tau_0\mu_0 + \tau\sum_{i=1}^{K_{tj}} x_i}{\tau_0 + \tau K_{tj}}, \frac{1}{\tau_0 + \tau K_{tj}}\right)$$

where $K_{tj} = \sum_{n=1}^{t} \mathbb{1}(L_{Gnj} = 1)$ and L_{Gnj} is an indicator variable equal to 1 if employer j hires a group G worker in period n.

Employers are wage-takers and maximize expected profits by hiring one worker each period

$$\begin{aligned} V(\psi_{S_{tj}}, w_{Gt}(\Psi_t)) = Max\{&\mu - w_O + \beta E_t[V(\psi_{S_{t+1,j}}, w_{Gt+1}(\Psi_{t+1}))], \\ &E_t[\mu_G|S_{tj}] - w_{Gt}(\Psi_t) + \beta E_t[V(\psi'_{S_{t+1,j}}, w_{G,t+1}(\Psi_{t+1}))]\} \end{aligned}$$

where $\psi_{S_{tj}} = \{E[\mu_G|S_{tj}], \mathrm{Var}(\mu_G|S_{tj})\}$ summarizes an employer's beliefs and Ψ_t is a list of beliefs across employers. When employers hire from group G, $V(\cdot)$ includes updated beliefs, and remains constant otherwise ($\psi_{S_{tj}} = \psi_{S_{t+1,j}}$). I focus on individual employer decision-making, taking worker wages as exogenous constants ($w_{Gt}(\Psi_t) = w_G = w_O$).[5] There is a trade off between benefits and costs of learning about group G, and whether employers hire from the group depends on their previous experiences. An employer hires from group G in period t if

$$\beta E_t[V(\psi'_{S_{t+1,j}}) - V(\psi_{S_{t+1,j}})] \geq \mu - E_t[\mu_G|S_{tj}]$$

which contrasts the expected information value of hiring from group G with the (potentially negative) expected cost in terms of foregone productivity. Optimal hiring follows a cutoff rule (Gittins and Jones, 1974). Employers hire from group G as long as their posterior beliefs are such that their expected value from hiring G is higher

(Anderson, 2001; 2012).

[5] Chapter 2 shows that endogenizing wages leads to a lower wage for group G.

than that of hiring O. If it starts below or falls below, employers switch to group O for all future periods since the decision problem remains unchanged.

Positive experiences, defined as hiring group G workers with productivity above that expected from group O, lead employers to continue hiring from group G on expected productivity grounds and learn about their productivity. Negative experiences, defined as hires with productivity below that expected from group O, lower the value of hiring from the group, reducing hiring and preserving beliefs. Negative biases are more persistent than positive ones, leading to a negatively-skewed belief distribution across employers.

To summarize, the mechanism leads to the following set of hypotheses:

- Hypothesis 1. Positive hiring experiences lead to a higher estimate of group G's mean productivity and more hiring from the group.
- Hypothesis 2. Negative hiring experiences lead to a lower estimate of group G's mean productivity and less hiring from the group.
- Hypothesis 3. Through increased hiring, positive experiences increase learning and lead to more accurate beliefs about group G's productivity.
- Hypothesis 4. Through decreased hiring, negative experiences decrease learning and lead to less accurate beliefs about group G's productivity.
- Hypothesis 5. Since negative biases are more persistent than positive ones, the final belief distribution about group G's productivity is negatively-skewed across employers.

3.3 Experimental Design

The experiment documents how an employer's hiring experiences with group G workers influence their subsequent hiring and beliefs about the group's productivity.

3.3.1 Workers

To construct a hiring pool for employers, workers were assigned the real-effort cognitive task of solving character puzzles under a piece rate of 250 credits. An example puzzle is shown in the Appendix B. Workers were given one practice puzzle followed by 4 minutes to solve as many puzzles as they could, which corresponds to their productivity in the experiment.

Based on the treatment, workers were randomly assigned to group *Gray* or *Orange*.[6] The two groups are equally productive by construction, but differ in relative size across employer treatment arms as described below.

3.3.2 Employers

Employers were incentivized with hiring the most productive workers over fifteen periods $t = 1, ..., 15$, which corresponded to hiring from the worker group with higher expected productivity, if any.[7] They observed their hire's productivity y_{it} each period and received 220 credits per puzzle solved by their hire for a total of $220y_{it}$, paid for a random subset H of 5 periods. After the instructions, employers answered comprehension questions to ensure a good understanding of the task.[8]

Before hiring, employers were informed about the worker task and shown an example puzzle. They were given the following initial information: the size of the worker groups and the mean productivity of group O, μ. When hiring from that group, they were given a worker with productivity equal to the group average, $y_{it} = \mu$, making concrete the idea that employers knew what to expect. Theoretically, this simplification has no impact on employer behavior based on expected productivity.

[6]To control for preferences, colors green and purple were also used and color order was varied such that some employers saw green or orange as the uncertain group and others purple or gray. Little difference was found based on colors and they are pooled together.

[7]Chapter 4 reports that the median number of hires by managers at a large US retailer is 10 over a six year period, suggesting that fifteen hires corresponds to a substantial real world time frame.

[8]Participants could attempt to answer the questions as many times as they wished within a one hour period, but could not continue or receive payment without answering all questions correctly. Some participants did not complete the questions and abandoned the experiment, likely improving data quality.

In practice, it simplifies the instructions substantially since the notion of averages is easier to grasp than distributions, expectations, and dispersion. It is also of little consequence for identifying the mechanism of interest, since I focus on the impact of hiring experiences on subsequent hiring and beliefs, rather than baseline hiring differentials across groups.[9] To directly investigate the role of ambiguity aversion, employers completed a separate task after their hiring task to obtain an individual measure of ambiguity aversion following Gneezy et al. (2015).

Beliefs were elicited using a binarized scoring rule as proposed in Hossain and Okui (2013) and incentivized for a random sample of two periods R.[10] It was made operational as follows: when employers hired from group O, their beliefs about group G carried over from the last period. Each period, current beliefs were used to compute a squared prediction error $(\mu - \mu_{Gjt})^2$. If the period was selected for payment, employers received 110 credits if their squared prediction error was below some number N_t and nothing otherwise. N_t was drawn each period from a uniform distribution on $[0, 81]$, with the upper limit selected to have a high probability of being larger than the squared prediction error under truthful reporting. Implicitly, employers learned about both the mean and the variance of group G productivity, but the belief elicitation mechanism isolates learning about the mean to focus on the impact of experiences on mean posterior beliefs. The total payoff of employer j corresponds to

$$\pi_j = \Sigma_{t=1}^{15} \mathbb{1}\{t \in H\} 220 y_{it} + \Sigma_{t=0}^{15} \mathbb{1}\{t \in R \cap (\mu - \mu_{Gjt})^2 < N_t\} 110.$$

3.3.3 Treatment Arms

To test hypotheses 1-5, employers were assigned to one of two treatments:

- Treatment B. Each period, employers choose between hiring from group O or

[9] The simplification gives risk-averse employers an incentive to hire from the certain group, but does not interact with the nature of hiring experiences whose impact is the object of interest. That is, the goal is not to document whether employers hire more or less workers from the uncertain group, but whether better or worse hiring experiences with the group cause relative changes in subsequent hiring and learning.

[10] The advantage over the quadratic scoring rule is that it does not require risk neutrality. It is truth-inducing as long as agents prefer lotteries with higher probabilities of larger payments.

G. Group O is framed as the majority group with 75% of workers. Beliefs about the mean productivity of Group G are elicited before the first hire, after every hire from the group, and at the end of the hiring task.

- Treatment C. As in Treatment B, but employers can only hire from Group G each period.

Treatment B represents the baseline task, providing a setting to test hypotheses 1 and 2 by observing how hiring experiences impact subsequent hiring and beliefs about group G's productivity.

Testing hypotheses 3 and 4 is complicated by the fact that hiring experiences affect posterior beliefs in two distinct ways: they mechanically lead to belief updating and may also impact hiring, indirectly affecting beliefs. The first results from standard belief updating, while the second corresponds to the mechanism of interest. For example, contrast the difference in mean posterior beliefs about group G's average productivity between information sets $S = s_1, ..., s_K$ and $S^{'}$ with $s_1^{'} < s_1$ and $s_i^{'} = s_i$ for $i = 2, ..., K$

$$E[\mu_G|S] - E[\mu_G|S^{'}] = \Delta_{\mu_G \mu_G^{'}}.$$

To measure the total effect of a lower first signal, contrast the difference in mean posterior beliefs between information sets S and $S^{''}$ with $J < K$ signals, $s_1^{''} = s_1^{'} < s_1$ and $s_i^{''} = s_i^{'} = s_i$ for $i = 2, ..., J$

$$E[\mu_G|S] - E[\mu_G|S^{''}] = \Delta_{\mu_G \mu_G^{''}}.$$

The impact of the mechanism of interest can be isolated by taking the difference

$$\Delta_{\mu_G^{'} \mu_G^{''}} = \Delta_{\mu_G \mu_G^{''}} - \Delta_{\mu_G \mu_G^{'}}.$$

Treatment C allows me to separately identify the two components by providing exogenous variation in beliefs that is uncorrelated with hiring experiences. For treatment C, hires influence beliefs about the productivity of Group G, but endogenous learning is shut down since employers cannot stop hiring from the group. Since treatment B employers generally hire from both groups regardless of their hiring experiences, I use beliefs of employers assigned to treatment C after the number of periods corresponding to the average number of G workers hired by treatment B when comparing between the two treatments.[11] Lastly, contrasting the final distributions of beliefs across the two treatments also allows me to test hypothesis 5.

To further investigate how the hiring context affects the formation of biased beliefs through endogenous employer learning about worker groups, I consider the following additional treatments:

- Treatment B1. As in Treatment B, but groups are framed as equally sized with 50 workers each.
- Treatment B2. As in Treatment B, but beliefs about the mean productivity of Group G are only elicited at the end of the hiring task.

Comparing Treatment B to Treatment $B1$ investigates how the framing of group G as a minority impacts employer behavior and beliefs, since minority status is a frequent feature of discrimination which may itself impact how employers form beliefs. Comparing Treatment B to Treatment $B2$ investigates how belief elicitation itself impacts behavior and beliefs. Belief elicitation incentives were chosen to be small compared to hiring payoffs in order to minimize distortions in hiring incentives. Still, eliciting beliefs could make employers more careful or suggest to them that their beliefs should change. Table 3.2 summarizes the different treatments.

3.4 Data

A group of 200 workers and 869 employers were recruited through Amazon's Mechanical Turk (MTurk), using an exchange rate of 1000 credits for \$0.2 and a subject pool restricted to US adults.[12] The experiment was implemented using oTree (Chen

[11] The main conclusions of the analysis are not sensitive to this choice of period.

[12] Data gathered through MTurk have been found to be reliable and consistent with data obtained from a traditional laboratory environment or other survey methods (Buhrmester et al., 2011; Berinsky et al., 2012; Goodman et al., 2013; Bentley et al., 2017; Mortensen and Hughes, 2018).

et al., 2016). Workers received a participation fee of \$0.75 in addition to their earnings, for an average total of \$1.25. Their study lasted approximately 7 minutes, corresponding to an hourly rate of around \$10-\$12. Summary statistics on workers are presented in Table 3.1. They solved 9 puzzles on average, with a minimum of 1 and a maximum of 18.

Employers received a participation fee of \$1 plus their earnings from the experiment, for a total of approximately \$3 on average. The study lasted around 12-15 minutes, corresponding to an hourly rate of around \$12-\$15.[13] The following sample restrictions were applied. A subset of employers (approximately 8%) who only hired from group O across all periods were excluded since they provide no usable variation.[14] Employers who reported beliefs higher than the maximum number of puzzles solved by workers (18) after the first hiring period were also excluded (approximately 2%) to avoid extreme beliefs influencing the analysis. In total, based on power calculations from pilot experiments and after applying sample restrictions, 281 employers were assigned to Treatment B, 139 to Treatment C, 182 to Treatment $B1$, and 185 to Treatment $B2$. Demographic information on employers is presented in Table 3.1. Slightly more than half are male, three-quarters are white, two-thirds have a college education, and three-quarters have some employment beyond MTurk.

3.5 Empirical Strategy

Hypotheses 1 and 2 posit that positive (negative) hiring experiences lead employers to update their beliefs about group G's mean productivity upwards (downwards), increasing (decreasing) subsequent hiring of the group. I investigate this by estimating the following models on Treatment B employers:

$$Y_{j,t+1} = \beta_0 + \beta_1 P_{jt} + \alpha_j + \varepsilon_{jt} \tag{3.1}$$

[13] Employers and workers were calibrated to earn the same hourly rate, but employers finished the task quicker than expected on average. Employers and workers were not made aware of each other's earnings.

[14] These employers did not report substantially lower priors about group G productivity or higher ambiguity aversion. Debriefing suggests that they used a simple strategy to quickly obtain a guaranteed payoff.

and

$$F_{15j} = \beta_0 + \beta_1 P_{jt} + \beta_2 X_j + \varepsilon_{jt} \tag{3.2}$$

where $Y_{j,t+1}$ corresponds to whether the next hire is from group G or to beliefs about group G's mean productivity carried into the next period. F_{15j} corresponds to the total number of hires from group G or to beliefs about their mean productivity after 15 periods. P_{jt} measures the productivity of previous group G hires. I consider several definitions, such as the mean productivity of group G workers over a range of periods and indicators for hires with productivity below or above 9 (the mean productivity of group O) in a subset of early periods. α_j is a collection of employer fixed effects included to capture time-invariant tendencies across employers to hire from a given group or update their beliefs.[15] X_j is an individual measure of ambiguity aversion. Lastly, standard errors are clustered at the employer level for equation (3.1).

Hypotheses 3 and 4 posit that positive (negative) hiring experiences increase (decrease) learning about group G by affecting hiring, leading to more (less) accurate beliefs about the group's productivity. The main outcome of interest is $|\mu_{Gtj} - \mu|$, a measure of biased beliefs about group G's mean productivity. The experimental data is used to estimate the following empirical model comparing treatments B and C

$$|\mu_{Gtj} - \mu| = \beta_0 + \beta_1 P_{jt} + \beta_2 B_j + \beta_3 P_{jt} * B_j + \varepsilon_{jt} \tag{3.3}$$

where B_j is an indicator for the employer having been assigned to treatment B. The coefficient of interest, β_3, represents the additional impact of hiring experiences for employers who can hire from either group, isolating the mechanism of interest.

[15] Results are similar without fixed effects and qualitatively similar using a first-difference estimator.

3.6 Results

I provide evidence for hypotheses 1 and 2 in Tables 3.3 and 3.4. Table 3.3 presents estimates of equation (3.1) for Treatment B. Estimates regarding hiring correspond to a 0.02 percentage point or 3% increase in the probability of hiring group G for each additional puzzle solved by the previous G hire and a 24% decrease (21% increase) if the previous worker was below (above) the mean Group O productivity of 9 puzzles. Estimates regarding beliefs correspond to a 2% or 0.17 puzzle increase in beliefs about the mean productivity of group G for each puzzle solved by the previous group G hire and a 12% or 1 puzzle decrease (increase) if the previous worker was below (above) the Group O average. Estimates are statistically significant at the 1% level. Experiences in the last period clearly influence current hiring and beliefs, and the impact appears relatively symmetric for positive and negative experiences.

Employers switch between groups on average 3.86 times with a standard deviation of 3.17. Still, 28% switch at most once and 43% switch at most twice, indicating that a substantial fraction of employers were at or close to the optimal number. Increased switching could mitigate the impact of the mechanism since employers do not completely stop hiring group G workers after switching to group O once. Yet, employers may be quicker to switch away from group G in the first place and since the impact of early learning may be particularly important, this may in turn decrease hiring and learning about group G.

Table 3.4 presents estimates of equation (3.2) for Treatment B. The impact of hiring experiences on total hiring can be seen as a "first stage", since the mechanism posits that hiring experiences impact beliefs specifically though changes in hiring. Estimates from the first three columns of the top panel show a strong statistically significant relationship between total hires and the productivity of the first, the average of the first three, and the average of all workers hired from group G (3%, 7% and 15% increases per additional puzzle solved).[16] The first six columns of the bottom panel show that early experiences with group G are an important determinant of total hiring from the group. The first three columns show that hiring a first, first two, or first

[16]The sample size changes across columns since employers who stop hiring from group G after their first (second) hire from the group are excluded when calculating the average of the first two (three) hires.

three workers below average productivity is associated with statistically significant lower hiring of the group (26%, 27% and 30% decreases). Contrastingly, the impact of positive early experiences progressively decreases from a 19% increase for the first hire to a non-statistically significant 8% increase for the first three hires. Larger magnitudes for negative experiences are expected because they lead to avoidance, while positive experiences lead to more hiring which mitigates their impact on average. Throughout, including the individual measure of ambiguity aversion has little impact on the magnitude or statistical significance of the estimates of interest.

Columns 7-12 of Table 3.4 show a strong statistically significant relationship between early hiring experiences and final beliefs about the group, but conflate the impact of experiences on beliefs through hiring and mechanical belief updating. Estimates from the top panel correspond to 1% , 5% and, 8% increases in final beliefs about group G's mean productivity for each additional puzzle solved by the first, the average of the first three, and the average of all G workers. Estimates from Columns 7-9 of the bottom panel grow larger with additional negative experiences, corresponding to decreases of 8% , 14%, and 26% in final beliefs if the first, first two, and first three hires from the group have below-mean productivity. The relationship with positive early experiences is weaker, as expected, corresponding to 6-13% increases in final beliefs. Overall, the evidence from Tables 3.3 and 3.4 strongly supports hypotheses 1 and 2.

Next, I provide evidence relating to hypotheses 3 and 4 shown in Table 3.5. First, column 1 of the top panel provides evidence for Treatment B that each group G hire decreases final bias by 5%. I then estimate equation (3.3) and present estimates of the interaction term between hiring experiences and employers being assigned to Treatment B. For Treatment C, beliefs after 9 periods are used to construct the measure of final bias, since Treatment B employers hired 9 G workers on average. Estimates from columns 2-4 of the top panel are statistically significant at the 5% level and provide evidence that bias falls by 6%, 15% and 17% more for Treatment B for a one puzzle increase in the productivity of the first, the average of the first three, and the average of all group G hires. Estimates in columns 1-3 of the bottom panel are statistically significant at the 5% level and correspond to additional increases in bias of 55%, 59%, and 103% if the first, first two, and first three group G hires have below-average productivity. Estimates in columns 4-6 regarding the impact of positive experiences correspond to additional decreases in bias of 51%, 27%, and 20%, although only the estimate in column 4 is statistically significant. Additional negative experiences have an increasing impact on bias for Treatment B, but the effect of positive experiences is decreasing since they increase hiring and learning so that both employer treatments observe a more similar number of signals. The evidence provides direct support for hypotheses 3 and 4 at the core of the endogenous bias-generating mechanism.

I provide evidence for hypothesis 5 in Figure 1, contrasting the change in the distribution of employer beliefs between Treatments B and C from the first period of hiring to the last. Both treatments have slightly negatively-biased initial beliefs about the mean productivity of group G, with an average of 8.6. Treatment C generally corrected their biases, with both tails of the distribution shrinking, increased mass around 9, and average beliefs of 9.05 after 15 periods. In contrast, Treatment B have essentially the same average beliefs as in period 1. While the right tail of the distribution shrunk, proportionally little changed in the left tail. The skewness in the final belief distribution is -0.33 for Treatment B versus 0.08 for Treatment C, and both Wilcoxon rank-sum and Kolmogorov-Smirnov tests reject the null hypothesis of equal distributions at the 10% level. This highlights how employer beliefs about group productivity may not converge or converge slowly with experience, when experience is itself endogenous to an employer's hiring history.

Table 3.6 contrasts the impact of hiring experiences on total hiring and final beliefs between Treatments B, $B1$ and $B2$. Overall, although interaction terms are generally not statistically significant and negative experiences decrease hiring and beliefs for both treatments, estimates fairly consistently suggest that the impact of negative experiences is larger when group G is framed as a minority, as indicated by the positive interaction terms with Treatment $B1$ in columns 1-3 and 7-9 of the bottom panel. Employers appear quicker to draw a conclusion regarding the relative productivity of group G when it is presented as having relatively few workers, consistent with stereotyping. Although interaction terms with treatment $B2$ are also generally not statistically significant, estimates from columns 1-3 and 7-9 of the bottom panel consistently suggest that the impact of negative experiences is larger when beliefs are only elicited at the end. This is consistent with belief elicitation throughout the hiring task making employers more careful in their evaluation of group G.

3.6.1 Employer Characteristics

After the hiring task, employers completed an exit survey asking them some demographic information and their views on some race-related questions from the General Social-Survey. Table 3.7 relates characteristics of employers to hiring of group G and bias about their mean productivity. Participants with a higher measure of prejudice based on their average answer to six race-related questions hired 15% fewer G workers when the group was presented as a minority.[17] The estimated coefficient is much smaller when groups were presented as equally-sized. In both cases, the relationship with final beliefs appears negligible. The interaction between prejudice and the nature of hiring experiences was also investigated, but revealed little beyond an additional impact of positive experiences on future hiring for more prejudiced employers. Overall, there is little evidence of systematic relationships between other employer characteristics and hiring or bias.

3.6.2 Deviations from Bayesian Updating

The mechanism's impact on hiring and beliefs in practice could be affected by stereotype formation (Allport, 1954) and the law of small numbers (Rabin, 2002), among other factors (Kahneman, 2003).[18] Another benefit of the experiment is therefore quantifying the net impact of the mechanism, accounting for deviations from Bayesian updating.

The variance in group G productivity is unknown to employers, but their posterior mean updating can suggests particular deviations from Bayesian updating. For every round in which an employer reports their beliefs, I calculate their implied $t = 0$ para-

[17] Participants reported how much they agree (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree) with the following statements. 1 - In general, African-Americans are as hard-working as whites. 2 - In general, African-Americans are as competent at their job as whites. 3 - In general, African-Americans are as intelligent as whites. 4 - You would object if a family member brought an African-American friend home for dinner. 5 - There should be laws against marriages between African-Americans and whites. 6 - You would vote for an African-American candidate for president if they were qualified. African-Americans were asked to answer questions 4 and 6 replacing African-American with white.

[18] Previous results in the literature indicate that the context can be an important determinant of belief updating and lead to both under or over updating (Tversky and Kahneman, 1971; Grether, 1980; Mobius et al., 2014; Coutts, 2019; Enke and Zimmermann, 2019).

meter κ_0, which represents initial beliefs about variance in productivity of workers.[19] Under Bayesian updating, κ_0 is a positive time-invariant constant, with a higher value implying lower initial beliefs about variance and therefore more updating conditional on a signal.

An increasing κ_0 across periods suggests potential over-updating, consistent with employers updating about the mean by more than implied from their initial beliefs about the variance. κ_0 can also be negative if posterior mean beliefs are above or below both μ_0 and $\bar{x}$, or equal to infinity if employers do not update at all. More precisely, a negative κ_0 is consistent with over-updating when employers update "too much" away from their prior towards $\bar{x}$. For example, this arises if an employer with prior beliefs of 9 observes signals of mean 8, but reports posterior beliefs of 7. Alternatively, a negative κ_0 is consistent with over-weighting of positive or negative experiences, such that prior beliefs are closer to $\bar{x}$ than posterior beliefs. For example, this arises if an employer with prior beliefs of 9 observes signals of mean 8, but reports posterior mean beliefs of 10.

Table 3.8 summarizes the implied values of κ_0 across employers, separating negative values based on whether they are consistent with over-updating or over-weighing. It also investigates whether these values change with experience or the productivity of the last group G worker hired. The table shows in columns 2-3 that the majority of non-missing values for κ_0 are negative, primarily consistent with over-updating rather than over-weighting, and becoming more frequent with hiring experience. Restricting to the subset of κ_0 with positive values, column 1 indicates that the magnitude of κ_0 increases with hiring experience, also consistent with over-updating. Around 26% of implied κ_0 values are missing, presumably arising from most employers reporting their beliefs as integers, and its frequency does not change with hiring experience as indicated by column 4. Across columns, there is little evidence that updating patterns vary with the productivity of the last hire. Overall, results are consistent with stereotype formation, where employers over-estimate the homogeneity of group G and update from relatively little information. This exacerbates the impact of the bias-generating mechanism given the different impact of positive versus negative

[19]The conjugate prior of a normal distribution with unknown mean and variance is the normal-gamma distribution. The closed form expression for the posterior mean corresponds to $\mu_n = \frac{\kappa_0\mu_0+n\bar{x}}{\kappa_0+n}$. From this expression, it is straightforward to recover the implied κ_0, given that everything else is observed.

experiences on future hiring and learning.

3.7 Conclusion

This paper provides evidence that employers endogenously develop negatively-biased beliefs about the productivity of a worker group about whose productivity they are initially more uncertain through their hiring experiences with its workers. Employer experiences with the group determine subsequent incentives to hire from the group again and learn about their productivity. I find that employers reduce their hiring of the group following negative experiences, decreasing their learning and leading to negatively-biased beliefs. In contrast, positive biases are relatively temporary since they endogenously lead to more hiring and learning. Across employers, since negative biases are more persistent than positive ones, this leads to a negatively-skewed distribution of beliefs.

I also show that the hiring context matters for the formation of biased beliefs. Namely, whether a worker group is framed as a minority and whether employers are primed to regularly report their beliefs both affect hiring and learning about group G. Lastly, I provide evidence that employers in practice are quicker to form negative stereotypes about worker groups than a Bayesian benchmark, implying that biased beliefs that arise through endogenous learning may be more pervasive than predicted.

The paper studies a fundamental feature of hiring in labor markets: in many settings, employers through their hiring learn not only about individual worker productivity, but also that of their group. Learning about groups is seldom considered in models of employer learning and discrimination, but I present evidence that it can play an important role in shaping employer behavior and worker outcomes. Importantly, biased beliefs in this paper do not rely on prior bias, prejudice, or productivity differentials across worker groups, providing evidence a new potential source of discrimination in the labor market.

Figure 3.1: Distribution of Employer Beliefs

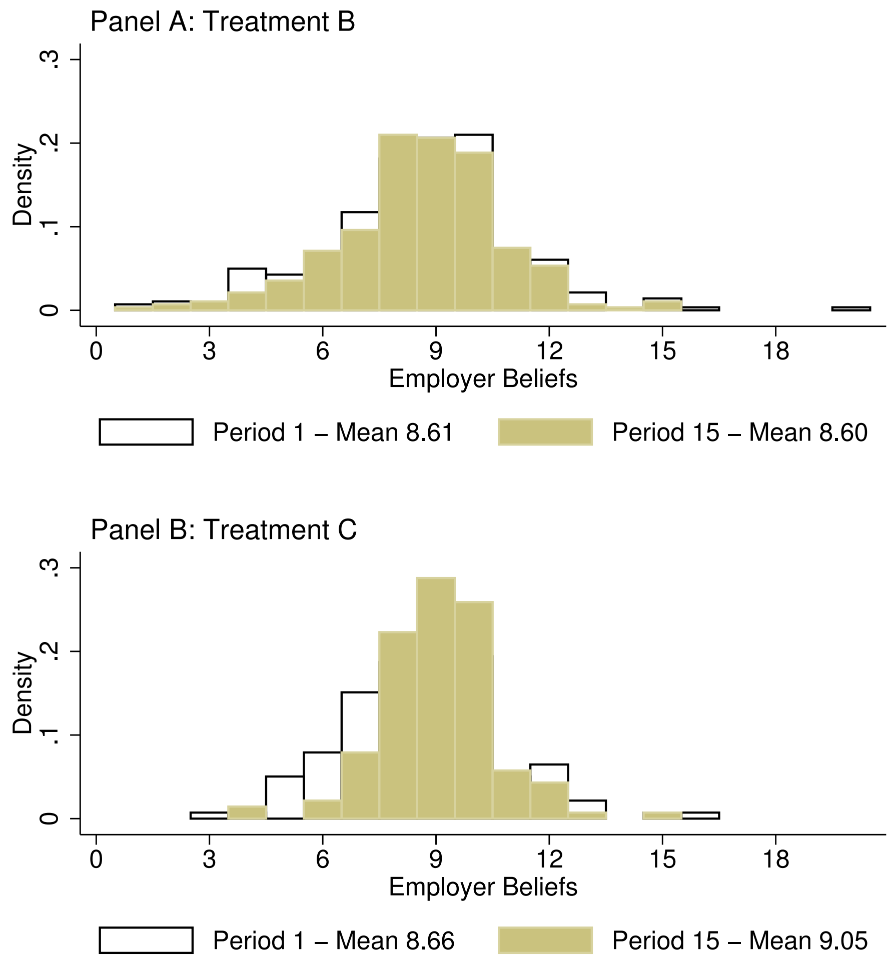

See Table 3.2 for a description of treatment groups.

Table 3.1: Summary Statistics

Panel A: Puzzles Solved by Workers

	Orange	Gray
Mean	9.23	9.12
Std. Dev.	3.44	3.68
Median	9	9
Min	1	1
Max	18	18
N. Obs.	150	50
P-value		
H_0: $\mu_O = \mu_G$	0.85	

Panel B: Employer Demographics

	Mean	Std. Dev.
Age	35.47	10.35
Male	0.58	0.49
White	0.75	0.43
Black	0.09	0.28
Asian	0.08	0.27
Hispanic	0.06	0.24
College Graduate	0.63	0.48
Employed Outside Mturk	0.73	0.45
N. Obs.	787	

Table 3.2: Employer Treatment Groups

	N	Hiring	Minority Framing	Belief Elicitation
Treatment B	281	Group G or O	Group G Minority	Prior, after every G hire, at the end
Treatment $B1$	182	Group G or O	Equal Group Sizes	Prior, after every G hire, at the end
Treatment $B2$	185	Group G or O	Group G Minority	At the end
Treatment C	139	Group G	Group G Minority	Prior, after every G hire, at the end

When Group G was framed as a minority group, it was presented as having 25% of workers (50). When both groups were framed as equally sized, they were presented as having 50 workers each.

Table 3.3: Impact of Previous Gray Hire on Current Hiring and Beliefs, Treatment B

	Prob. of Hiring Gray			Current Beliefs		
Prev. Worker Prod.	0.020			0.172		
	(0.002)			(0.011)		
Prev. Worker > 9		0.124			1.038	
		(0.017)			(0.068)	
Prev. Worker < 9			-0.141			-1.001
			(0.019)			(0.070)
Outcome Mean	0.59	0.59	0.59	8.60	8.60	8.60
N. Obs.	2,306	2,306	2,306	2,465	2,465	2,465

Clustered standard errors at the employer level are presented in parentheses. Regressions include employer fixed effects. Prev. Worker Prod. refers to the productivity of the worker hired from group G last period. Prev. Worker > 9 ($<$ 9) refers to whether the worker hired from group G last period had productivity above (below) the mean productivity of group O, 9. The number of observations differs across outcomes since a group G worker hired in period 15 leads to belief updating but no further hiring. See Table 3.2 for a description of treatment groups.

Table 3.4: Impact of Hiring Experiences with Gray Workers on Hiring and Final Beliefs, Treatment B

	Total Gray Hires						Final Beliefs					
Worker #1 Prod.	0.277						0.088					
	(0.076)						(0.034)					
Worker #1-3 Avg. Prod.		0.659						0.409				
		(0.128)						(0.068)				
Worker Avg. Prod.			1.334						0.723			
			(0.145)						(0.102)			
Outcome Mean	8.77	9.72	8.77				8.60	8.76	8.60			
N. Obs.	281	249	281				281	249	281			
Worker #1 < 9	-2.248						-0.657					
	(0.532)						(0.255)					
Worker #1-2 < 9		-2.505						-1.225				
		(0.614)						(0.341)				
Worker #1-3 < 9			-2.939						-2.246			
			(0.963)						(0.616)			
Worker #1 > 9				1.672						0.516		
				(0.540)						(0.253)		
Worker #1-2 > 9					1.499						1.110	
					(0.588)						(0.303)	
Worker #1-3 > 9						0.770						1.097
						(0.902)						(0.461)
Outcome Mean	8.77	9.34	9.72	8.77	9.34	9.72	8.60	8.68	8.76	8.60	8.68	8.76
N. Obs.	281	262	249	281	262	249	281	262	249	281	262	249

Robust standard errors are presented in parentheses. Regressions include an individual measure of ambiguity aversion calculated as in Gneezy et al. (2015). Employers who stop hiring from group G after the first (two first) hire(s) are excluded when calculating the average productivity of the first two (three) G workers because these employers have less than two (three) total experiences with the group. See Table 3.2 for a description of treatment groups and Table 3.3 for definitions.

Table 3.5: Differential Impact of Hiring Experiences with Gray Workers on Final Bias, Treatment B versus C

	Final Bias					
Total Gray Hires	-0.086					
	(0.019)					
C*Worker #1 Prod.		-0.082				
		(0.034)				
C*Worker #1-3 Avg. Prod.			-0.198			
			(0.068)			
C*Worker Avg. Prod.				-0.242		
				(0.121)		
Outcome Mean	1.62	1.45	1.35	1.45		
N. Obs.	281	420	388	420		
C*Worker #1 $<$ 9	0.798					
	(0.275)					
C*Worker #1-2 $<$ 9		0.810				
		(0.333)				
C*Worker #1-3 $<$ 9			1.416			
			(0.587)			
C*Worker #1 $>$ 9				-0.700		
				(0.283)		
C*Worker #1-2 $>$ 9					-0.365	
					(0.280)	
C*Worker #1-3 $>$ 9						-0.270
						(0.378)
Outcome Mean	1.45	1.37	1.35	1.38	1.37	1.35
N. Obs.	420	401	388	420	401	388

Robust standard errors presented in parentheses. Final bias is defined as the absolute value of the difference between employer beliefs and the true mean group productivity. Beliefs in period 9 are used to construct the bias measure for Treatment C and beliefs in period 15 are used to construct the bias measure for Treatment B. Treatment B employers who stop hiring from group G after the first (two first) hire(s) are excluded when calculating the average productivity of the first two (three) G workers because these employers have less than two (three) total experiences with the group. See Table 3.2 for a description of treatment groups and Table 3.3 for definitions.

Table 3.6: Differential Impact of Hiring Experiences on Total Hiring and Final Beliefs, Treatments B versus Treatments B1 and B2

	Total Gray Hires						Final Beliefs					
Worker 1 Prod.	0.279						0.091					
	(0.076)						(0.034)					
B1*Worker #1 Prod.	-0.098						0.018					
	(0.109)						(0.051)					
B2*Worker #1 Prod.	0.148						0.277					
	(0.108)						(0.200)					
Worker #1-3 Prod.		0.659						0.410				
		(0.128)						(0.069)				
B1*Worker #1-3 Prod.		-0.288						-0.169				
		(0.200)						(0.105)				
B2*Worker #1-3 Prod.		0.245						0.301				
		(0.187)						(0.267)				
Worker Avg. Prod.			1.336						0.726			
			(0.145)						(0.102)			
B1*Worker Avg. Prod.			-0.282						0.072			
			(0.283)						(0.147)			
B2*Worker Avg. Prod.			0.015						0.153			
			(0.231)						(0.153)			
Outcome Mean	8.96	9.78	8.96				8.85	9.01	8.85			
N. Obs.	648	584	648				648	584	648			
Worker #1 < 9	-2.259						-0.676					
	(0.532)						(0.256)					
B1*Worker #1 < 9	0.428						0.152					
	(0.836)						(0.389)					
B2*Worker #1 < 9	-0.206						-1.361					
	(0.806)						(1.069)					
Worker #1-2 < 9		-2.490						-1.215				
		(0.611)						(0.343)				
B1*Worker #1-2 < 9		1.357						0.323				
		(0.991)						(0.487)				
B2*Worker #1-2 < 9		-0.367						-0.676				
		(0.937)						(0.866)				
Worker #1-3 < 9			-2.956						-2.260			
			(0.952)						(0.624)			
B1*Worker #1-3 < 9			1.733						0.467			
			(1.480)						(0.737)			
B2*Worker #1-3 < 9			-1.428						-0.585			
			(1.400)						(1.059)			
Worker #1 > 9				1.677						0.524		
				(0.540)						(0.254)		
B1*Worker #1 > 9				0.276						0.008		
				(0.821)						(0.382)		
B2*Worker #1 > 9				0.376						1.549		
				(0.810)						(1.211)		
Worker #1-2 > 9					1.491						1.105	
					(0.588)						(0.302)	
B1*Worker #1-2 > 9					-0.086						-0.549	
					(0.888)						(0.483)	
B2*Worker #1-2 > 9					1.587						1.261	
					(0.835)						(2.258)	
Worker #1-3 > 9						0.744						1.074
						(0.905)						(0.468)
B1*Worker #1-3 > 9						-0.653						-1.108
						(1.187)						(0.666)
B2*Worker #1-3 > 9						2.776						-1.068
						(1.160)						(0.897)
Outcome Mean	8.96	9.38	9.78	8.96	9.38	9.78	8.85	8.91	9.01	8.85	8.91	9.01
N. Obs.	648	615	584	648	615	584	648	615	584	648	615	584

Robust standard errors presented in parentheses. See Table 3.2 for a description of treatment groups and Table 3.3-3.4 for definitions.

Table 3.7: Employer Characteristics, Hiring and Bias

	Total Gray Hires		Final Bias	
	Treatment B	Treatment B1	Treatment B	Treatment B1
Prejudice	-1.308	-0.550	0.046	0.229
	(0.423)	(0.543)	(0.119)	(0.169)
Less than college	-0.803	0.178	-0.058	-0.258
	(0.773)	(1.085)	(0.245)	(0.277)
Age	0.032	0.025	-0.002	0.019
	(0.026)	(0.026)	(0.008)	(0.010)
Male	0.196	0.220	-0.249	-0.247
	(0.577)	(0.650)	(0.197)	(0.204)
Employed	-0.512	0.025	-0.013	-0.040
	(0.594)	(0.741)	(0.213)	(0.241)
Black	-0.080	-1.649	0.685	0.040
	(0.813)	(1.213)	(0.386)	(0.295)
Hispanic	0.572	-2.812	0.351	0.033
	(1.012)	(1.309)	(0.364)	(0.385)
Outcome Mean	8.77	9.37	1.62	1.47
N. Obs.	281	182	281	182

Robust standard errors are presented in parentheses. Prejudice refers to an index measure based on average responses to six race-related questions adapted from the General Social Survey. Employed is an indicator variable for whether the participant is employed beyond their work on Mechanical Turk. See Table 3.2 for a description of treatment groups and Table 3.5 for definitions.

Table 3.8: Departures from Bayesian Updating, Treatment B

	κ_0 >0	κ_0 <0		
	κ_0	Over-Updating	Over-Weighing	Prob. κ_0 Missing
Number of Hires	0.470	0.035	0.010	0.002
	(0.115)	(0.005)	(0.003)	(0.004)
Worker Prod.	0.024	0.001	-0.003	-0.003
	(0.046)	(0.004)	(0.003)	(0.003)
Mean	3.212	0.389	0.100	0.260
N. Obs.	675	1,814	1,814	2,465

Clustered standard errors at the employer level are presented in parentheses. Regressions include employer fixed effects. See Table 3.2 for a description of treatment groups.

CHAPTER IV

The Emergence of Hiring Discrimination

joint work with Alan Benson

4.1 Introduction

A substantial body of research has examined the extensive and persistent racial disparities in the labor market, particularly studying the root causes of employment discrimination which is essential to develop theory, interpret evidence, and design policy interventions (Lang and Lehmann, 2012). One prominent strain of recent research has studied how idiosyncratic, biased beliefs among individual employers, judges, and other influential decision-makers can be enormously consequential in creating outcome differentials across groups (Reuben et al., 2014; Arnold et al., 2018; Bohren et al., 2019a; 2019b; Bordalo et al., 2019; Sarsons, 2019).

But why are the beliefs of managers and other decision makers biased? Previous work has proposed that biases may arise from implicit group associations (Bertrand et al., 2005) or the exaggeration of small differences in true group means (Bordalo et al., 2016). Chapter II proposes that such biases may also arise through market interactions with groups. If managers update their beliefs based on their previous experiences with a group, then these experiences also shape subsequent incentives to hire from the group and learn more about its workers. For example, early negative experiences may prompt managers to avoid hiring members of a certain group, which may itself slow learning. However, data limitations have stymied efforts to study the emergence of individuals' biased beliefs and discriminatory hiring in labor markets (Charles and Guryan, 2011; Guryan and Charles, 2013). For instance, Census and audit study data that have been workhorses of discrimination research rarely feature data on individual hiring managers or hiring within managers over time, inhibiting researchers' ability to study the emergence of biases.

In this paper, we use administrative data from the US operations of a large national retailer to examine how managers' past experience hiring workers of different races affects the race of their subsequent hires. The data, which include over 1 million permanent workers working under 27,000 department store managers across 4,900 stores between 2009 and 2016, 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 incentivized to hire a productive team and free to use any information gleaned from interviews or past experience to make decisions. The data 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 (BLS, 2021).

We begin by establishing substantial cross-manager heterogeneity in black hires, even after controlling for manager race, store, department, and job effects. Examining managers who move across stores, we find that a substantial share of the residual variation in black hiring is explained by manager fixed effects, implying that individual manager idiosyncrasies play a substantial role in determining the race of hires.

To examine whether variation in the race of hires across managers can be explained by biased beliefs seeded from their previous hiring experiences, we begin from a theoretical framework adapted from Chapter II. Managers are initially uncertain about differences in performance among applicants of different groups (e.g. white and black), but update their beliefs as they hire workers. Because hiring is based on managers' beliefs, and those beliefs also depend on hiring, learning is endogenous. Learning is also asymmetric: positive hiring experiences with a worker group lead managers to update their beliefs of that group's performance distribution upward, increasing their propensity to hire from that group. In contrast, negative experiences discourage future hiring, which also slows learning. Therefore, positive biases self-correct more quickly than negative ones. The model also yields the prediction that managers update their beliefs more following experiences with minority groups with whom they have less experience.[1] Combined with the relative persistence of negative biases, this systematically decreases the hiring of minority workers.

To operationalize positive and negative hiring experiences, we use dismissal and quit rates among workers hired for permanent positions. Turnover at this firm (and in retail generally) is very high, as are the cumulative costs of recruiting, training, and ramping up new workers. We define a positive experience hiring a worker of a given race as a spell in which turnover for that race is lower than expected given observable worker and job characteristics. We define negative experiences as higher-than-expected turnover. Our main results follow.

First, positive past experiences with white or black workers increase a manager's propensity to hire from that group, whereas negative experiences decrease their propensity. This result is consistent with the proposition that managers update their beliefs based on their personal experiences on the job, and not with the proposition that discriminatory beliefs and behaviors are stationary or "fixed" by the time they become managers.

Second, we find that learning is asymmetric by race. Although past experience affects subsequent hiring for both black and white workers, the effect is particularly pronounced for black workers. This suggests that managers have relatively weak priors about minority groups, and as a result, similar information yields greater changes in beliefs.

Third, early negative hiring experiences lead to a substantial persistent decline in black hiring; the effects of early experiences with white workers or early positive experiences with black workers are comparatively small and short-lived. The importance of early negative experiences with black workers suggests that the initial "seeding" of hiring experiences with minorities can have substantial, persistent effects on the

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

bias and hiring of individual managers.

We investigate several potential mechanisms through which managers' hiring histories could affect subsequent hiring, including supply-side responses by workers, hiring through referrals, and selection on unobservables. We conclude that these alternative mechanisms are unlikely to explain our findings. Rather, our set of results are most consistent with the proposition that managers update their beliefs about groups based on their own personal experiences, which in turn affects their future hiring.

These results have several implications for the study of discrimination. First, we provide evidence that race is a salient worker characteristic to managers and biased hiring occurs at the level of the manager based on their personal past experience on the job. Beliefs are not static, they are not purely determined at the level of the firm, and not fully checked by firm-level learning or signals provided by algorithmic recommendations and interviews. In contrast, much of the existing work focuses on time-invariant firm or manager effects, sometimes explained by variation in manager race (Giuliano et al., 2009; Giuliano and Ransom, 2013; Benson, Board and Meyer-ter-Vehn, 2019),[2] or implicit bias (Glover et al., 2017). Second, we provide evidence that hiring and learning are endogenously driven by managers' personal experiences with worker groups, therefore learning can be particularly slow for managers who have relatively little experience hiring minorities. Although we focus on hiring, this could presumably be extended to wages, as conventionally studied in employer learning models (e.g. Farber and Gibbons, 1996; Altonji and Pierret, 2001; Lange, 2007; Arcidiacono et al. 2010; Kahn and Lange, 2014; Lesner, 2018).

More generally, the endogenous formation of biased beliefs among employers and other economic gatekeepers is absent from much of the discrimination literature. We provide rare evidence consistent with fairly persistent biased beliefs in a broadly representative labor market setting and more particularly with endogenous learning about worker groups generating biased beliefs and discrimination. Making the distinction between endogenous biased beliefs and standard theories of statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977; Coate and Loury,

[2]Focusing on differences across manager racial groups can complicate interpretation, since they likely differ across unobservable characteristics which affect relative hiring, such as access to a network of workers from their own racial group.

1993), in which discrimination arises as a response of employers to true group differences, or taste-based discrimination (Becker, 1957), is crucial because they can lead to very different conclusions about how discrimination arises and can be mitigated. Our findings highlight that, even if employer beliefs about worker groups eventually converge with experience, algorithmic decision-making, or market competition, they do so slowly and unevenly across groups in a manner that can create discrimination. They also suggest that theories of statistical discrimination that allow for biased employer beliefs may capture important aspects of discrimination and help explain its pervasiveness and persistence.

4.2 Theoretical Framework

We present a model adapting Chapter II in which managers iteratively update their beliefs about the productivity of worker groups based on the observed productivity of their hires, and hire based on their (endogenous) beliefs. The main difference with the model in Chapter II is that we consider individual hiring decisions of managers who hire from an exogenous set of applicants for a given position. Our exposition considers managers who condition beliefs on race, but in principle, the model could be applied to any group characteristic, such as gender, education, or personality. Moreover, although managers in the model update their beliefs about the general productivity of black and white workers, we emphasize that managers may be only updating their beliefs as though past experiences are indicative of the potential quality of hires given their location or job. Although it's unclear whether managers generalize their experiences to the full universe of people of a given group, our empirics suggest that managers perceive their prior experiences as informative at least within the same parent company.

Our primary proposition is that employer beliefs about a worker group's productivity become biased because they depend on observed hires which themselves depend on beliefs. Managers "learn" from a biased sample of worker productivity, and this bias is seeded by good or bad previous experiences with workers from the group. Moreover, if managers have noisier information on the productivity of minority groups (Lundberg and Startz, 1983; Lang, 1986; Cornell and Welch, 1996; Morgan and

Várdy, 2009), then experiences with these workers play a disproportionate role in shaping subsequent hiring. The key novel predictions of our framework are as follows. First, negative hiring experiences decrease subsequent hiring of that group, and positive experiences increase subsequent hiring. Second, because the rate of hiring drives the speed of learning, negative biases will be more persistent than positive ones. Third, persistent negative biases from hiring experiences disproportionately impact workers from minority groups, both because managers come in with weak priors and because the rate at which minorities are hired is naturally low. These predictions do not rely on biased updating or on different worker groups having different productivity distributions: whatever the true productivity of minority groups is, employers will underestimate it.

As in standard models of statistical discrimination, managers hire workers based on their expected productivity and, in the absence of perfect information on individuals at the hiring stage, use group membership as a potentially relevant indicator of individual productivity. Unlike standard models of statistical discrimination, we do not assume that employers have complete information on group productivity or that their beliefs are confirmed in equilibrium, for example through endogenous worker responses. Rather, employers are initially uncertain about the relative productivity of groups and update their beliefs through hiring.[3]

Specifically, a manager is tasked with hiring the most productive worker from a candidate pool, taking vacancies and entry wages for a position as given. Managers predict the productivity of worker i from group g, x_{ig}, based on a noisy signal of individual productivity s_{ig}, which they observe prior to hiring, and group membership g. The individual signal is composed of the worker's productivity and an unbiased noise component, such that $s_{ig} = x_{ig} + \varepsilon_{ig}$ with $\varepsilon_{ig} \sim N(0, \sigma^2_{\varepsilon g})$. For example, it could include information from a candidate's resume, a pre-employment test, or an interview.

We consider two worker groups denoted by W and B. Worker productivity is nor-

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

mally distributed with mean μ_g and variance σ_g^2, such that $x_{ig} \sim N(\mu_g, \sigma_g^2)$. The productivity distributions of groups are static, but beliefs about these distributions are heterogeneous and evolve over time. For expositional simplicity, we assume that managers know the variance in group productivity σ_g^2 and the noisiness of individual signals $\sigma_{\varepsilon g}^2$ for $g \in \{W, B\}$, which both potentially vary across groups.[4] As such, we focus on posterior beliefs about the mean productivity of group g, $\hat{\mu}_g$, and define $S_{gmt} = \{x_{ign} : i \text{ from } g \text{ is hired by } m \text{ at time } n\}_{n=1}^{t}$ as the information set about workers from group g available to manager m at time t.

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

$$P_{igmt} = E[x_{ig}|s_{ig}, E[\hat{\mu}_g|S_{gm,t-1}]] = \gamma_{gmt}s_{ig} + (1 - \gamma_{gmt})E[\hat{\mu}_g|S_{gm,t-1}]$$

where $\gamma_{gmt} = \frac{\sigma_g^2 + Var[\hat{\mu}_g|S_{gm,t-1}]}{\sigma_g^2 + Var[\hat{\mu}_g|S_{gm,t-1}] + \sigma_{\varepsilon g}^2}$.[5] The manager hires the worker with the highest expected productivity out of the set of applicants A_t, with fraction F_{gt} from group g.[6] That is, worker i from group g is hired at time t if $P_{igmt} > P_{i'g'mt}$ for all $i' \in A_t$ from group g', and for $g' \in \{W, B\}$.

When manager m begins hiring, they have a prior belief distribution about group g's mean productivity $\hat{\mu}_g|S_{gm0}$. If they don't hire from group g at time t, they don't update their beliefs and $\hat{\mu}_g|S_{gm,t-1} \stackrel{d}{=} \hat{\mu}_g|S_{gmt}$.[7] If they hire from group g at time t, managers observe their hire's productivity x_{ig}, which is not randomly drawn because managers select workers based on expected productivity.

Accordingly, managers first form an expectation about x_{ig} given that worker i has

[4]Employers learning about productivity variance or individual signal precision through their experiences would affect the relative weight attached to individual signals versus group membership across worker groups, but not the substantive implication that positive (negative) experiences increase (decrease) a group's expected productivity, and therefore the probability of hiring from the group.

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

[6]As modeled in Chapter II, managers should also value learning about groups to improve subsequent hiring. We abstract from this feature since it does not change the direction of belief updating and that of its impact on hiring and learning across groups.

[7]We assume that information from a manager's own hire is particularly salient compared to signals of workers not hired or correlated updating across groups, abstracting from these sources of updating for shorter-run hiring decisions.

the highest expected productivity out of all applicants, given by

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

Second, managers update their beliefs from $\hat{\mu}_g|S_{gm,t-1}$ to $\hat{\mu}_g|S_{gmt}$. Posterior variance monotonically decreases with additional hires, while the direction of posterior mean updating depends on the discrepancy between expected and observed productivity of the hire

$$E[x_{ig}|P_{igmt} > P_{i'g'mt}\ \forall i' \in A_t, g' \in \{W, B\}] - x_{ig}. \tag{4.1}$$

If realized productivity is above or below expectation, denoted as a positive and negative hiring experience respectively, managers update their beliefs based on equation (4.1) with upwards updating when the difference is negative and downwards updating otherwise. A positive (negative) experience with group g increases (decreases) $E[\hat{\mu}_g|S_{gmt}]$ and $P_{igm,t+1}$ relative to group g', increasing (decreasing) the probability that a group g worker is hired in period $t+1$. In our data, we have performance measures but not whether they were above or below the manager's expectation, so we use different performance measures relative to other workers at the firm, which should inform expectations. These predictions don't rely on prior bias or prejudice and don't depend on Bayesian updating, potentially including a wide class of belief updating rules. The subjective assessment of an employer as to what constitutes a positive or negative experience may itself be biased and vary across groups, but it is precisely the manager's perception that is of interest, rather than the worker's objective performance.[8]

Contrasting the impact of positive and negative hiring experiences, a positive experience, through increasing the probability of subsequent hiring from the group, also increases the probability of observing signals about the group's productivity at time $t+1$ and beyond. Depending on the manager's prior, these additional signals may partially undo the impact of the positive experience through reversion to the

[8]While evidence on the impact of experiences with workers more likely to have been exogenously assigned to managers is presented in Appendix C, we purposefully focus on a manager's own endogenous hiring decisions since they are likely more salient to the manager and more likely to lead to updating about worker groups given that the manager should have a clearer expectation of performance from being in charge of assessing and selecting workers.

mean and lead to more accurate beliefs. In contrast, a negative experience, through decreasing the probability of subsequent hiring from the group, also decreases the probability of observing signals about the group's productivity at time $t+1$ and beyond. Regardless of prior beliefs, the impact may be more persistent, because it decreases learning. The relative size of worker groups may also interact with these impacts. If F_g is small, workers from the group are infrequently hired. Belief updating from experience is less likely to be pivotal for subsequent hiring, but may have more persistent impacts since learning is slower. Similarly, it may be easier (harder) for managers to avoid hiring (seek out) groups who constitute a minority of applicants following negative (positive) experiences.

Lastly, belief updating across groups may differ based on differences in employer priors, worker productivity distributions, and group size. Managers may have higher prior precision about group W's productivity because they are the majority group or because they are themselves primarily from group W. Similarly, if there are fewer workers from group B, managers should update more about the group given equivalent signals. Differential updating across groups, combined with more persistent impacts of negative experiences on subsequent hiring, predicts that hiring experiences systematically decrease relative group B hiring across managers.

To summarize, three main predictions follow from the theoretical framework:

- 1. More negative (positive) hiring experiences of a manager with workers of a given group decrease (increase) the manager's subsequent hiring of the group.
- 2. Negative experiences have more persistent impacts, because they decrease subsequent learning about the group's productivity.
- 3. Experiences of a manager with groups disproportionately affect minority hiring.

4.3 Data, Hiring, and Empirical Design

Our data consist of longitudinal administrative records on workers and managers from the US operations of a large national retailer between February 2009 and October 2016. For each worker and manager, we observe tenure, demographics, job, department, and location. We also observe employment termination including dismissals, quits, and layoffs along with stated reasons for dismissals and quits. Each store is led by one store manager and a set of department managers who hire for their respective department, allowing us to study hiring decisions of each department manager over time. We restrict our sample to workers hired into permanent 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). We focus on white and black workers because these are the two largest racial categories in our data, which make it most feasible to estimate managers' evolving hiring behavior. Hispanics are treated as a separate category in the data and are presented in Appendix C. Summary statistics on workers and managers are presented in Table 4.1.

When a department has a vacant position, managers have access to a pool of existing applicants and can post to recruit additional candidates. Applicants take an online screening test that classifies them into three recommendation tiers. Department managers observe screening test results, but are free to deviate from the algorithmic recommendation when selecting applicants to interview. Department managers are trained in behavioral interview techniques. For instance, a customer service applicant may be asked how they might respond to a hypothetical scenario, like dealing with a difficult customer, or to discuss an instance they confronted a problem and were proud of the solution they offered. New hires complete one week of online training and one week of shadowing an experienced worker before starting in their regular position.

4.3.1 Characterizing Black Hiring Across Managers

We begin by characterizing heterogeneity in the hiring of black workers across managers and examining how much of it is due to idiosyncratic variation across individual managers versus external factors. Although descriptive, this type of decomposition can rarely be done due to data requirements and is valuable because theories of discrimination differ fundamentally in the predicted role of individual managers. Under classical statistical discrimination, managers discriminate similarly around the true productivity distribution of each group; they are not individually and idiosyncratically biased. In stark contrast, idiosyncratic prejudice or bias are at the center of taste-based and belief-based discrimination.

Figure 4.1 shows the share of black and white workers hired over our sample by each manager. There is substantial heterogeneity in group hiring across managers. The mean share of black workers hired by managers is 20%, the median is 8%, and more than a quarter hire no black workers. The mean share of white workers hired is 56% and the median is 59%.

Many factors presumably contribute to heterogeneity in hiring across worker groups, such as store location. To estimate how much heterogeneity in black hiring is explained by manager effects net of other factors that may vary by store, department, job, time period, or economic condition, we take Abowd et al. (1999)'s approach of analyzing connected sets of workers, specifically managers who work in multiple stores.[9] 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

$$Black_{imjlt} = X_{mjlt}\beta + \gamma_m + \alpha_j + \lambda_l + \theta_t + \varepsilon_{imjlt} \tag{4.2}$$

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.[10] 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

[9] 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. (2015) and Benson et al. (2019).

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

recruiting in jobs, locations, periods, and market conditions associated with more black hires.

Figure 4.2 contrasts the predicted black hiring shares across managers to the actual values. By construction, predicted shares approximate the middle of the distribution. Especially without manager fixed effects, they fail to capture much of the bottom of the distribution, predicting that too many managers hire 10-30% black workers and too few hire less. Beyond manager fixed effects, the majority of the explanatory power comes from the store fixed effects, which capture store-level and area-level characteristics. Without manager fixed effects, the adjusted R-squared with store fixed effects alone is approximately 0.285, while that of the full model is approximately 0.3. Manager fixed effects alone explain 4-5% of the total variation in black hiring and roughly a third of the discrepancy between actual shares and those predicted by the model without manager fixed effects. Qualitatively, the model with manager fixed effects still under-predicts the share of managers who hire very few or no black workers, but the discrepancy is substantively smaller. This exercise suggests that, beyond store and contextual factors, the specific identify of the hiring manager is an important predictor of black hiring in a department. Appendix C presents results restricted to managers who hire at least 5 workers over our sample period, which has little substantive impact on the results, and analogous results for white hiring.

The distribution of manager fixed effects is plotted in Figure 4.3, with the right panel restricted to managers who hire at least 5 workers. 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 in the left panel appears fairly symmetric, while that in the right panel displays a slight negative skew. As shown in Appendix C, 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).

4.3.2 Empirical Design

We next turn to our main empirical analysis regarding how managers' experiences with black and white workers shape their subsequent hiring behavior. We organize the data into a manager-level panel with each observation corresponding to a month in which a manager hires at least one black or white worker, which we refer to as a hiring event. Our baseline analysis restricts our sample to managers who began hiring at the firm for the first time during our sample period. On average, managers hire workers approximately every two and a half months, totaling 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).

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

$$F_{gemlt} = \beta_1 E\bar{X}P_{g,e-1} + \beta_2 E\bar{X}P_{g',e-1} + X_{gemlt}\zeta + \theta_t + \lambda_l + \gamma_m + \varepsilon_{gemlt} \tag{4.3}$$

where the dependent variable is the share of group g workers hired in hiring event e by manager m in location l at time t, corresponding to the share of black hires in our main analysis. The primary coefficient of interest is β_1, capturing the impact of more positive or negative average hiring experiences with black workers up to hiring event $e-1$. Similarly, $E\bar{X}P_{g',e-1}$ captures more positive or negative average hiring experiences with group g' up to hiring event $e-1$, corresponding to hiring experiences with white workers in our main analysis. X_{gemlt} includes the fraction of full-time workers, fraction female, average age, total number of hires, number of previous hiring events, time since last hiring event, yearly state unemployment, and yearly state college attainment. θ_t, λ_l, and γ_m represent month and year, store, and manager fixed effects. Store fixed effects account for differences between applicant pools, local markets, and store-level characteristics faced by the manager, among other factors. Manager fixed effects account for fundamental manager differences, for example in ability hiring and managing different worker groups or in taste-based discrimination. Standard errors are clustered at the manager level.

This specification can be used to test our theoretical framework's three key predictions by considering different measures for $E\bar{X}P_{g,e-1}$ and $E\bar{X}P_{g',e-1}$. First, the impact of the cumulative average experience with a worker race in previous hiring events on current hiring of the race provides evidence for prediction 1, that is whether more negative (positive) experiences of a manager with a race impact their subsequent hiring of the race. Second, considering average experience over a subset of early hiring events investigates the (relative) persistence of different hiring experiences on subsequent hiring, which we posit operates by endogenously affecting employer learning about the racial group's performance. Third, comparing the impact of previous hiring experiences with black and white workers on subsequent relative hiring investigates whether previous experiences have larger impacts on minority hiring.

Our empirical analysis considers hiring decisions as a function of idiosyncratically positive or negative experiences given the true expected mean of each race. This requires us to distinguish each hiring event as either positive or negative versus expectations for both worker races. To calculate idiosyncratic, manager-level deviations in observed experience versus expectation, we first compute deviations in monthly turnover rates by race at the level of the manager's subordinates from expected turnover rates at the firm.[11] The cumulative average of these monthly deviations indicates how previous hires from each racial group were more or less likely to achieve a given tenure than expected.

We also consider a second approach separating particularly negative and positive experiences of a manager, focusing on specific experiences likely to be most salient to the manager. For negative experiences, for each hiring event, we calculate the share of each race that was fired or quit in the first 3 months of employment. Workers hired into permanent positions who leave or are terminated within the first 3 months account for around a quarter of hires and they are costly in terms of direct costs of hiring and training, foregone on-the-job skills training and ramp up period of a potentially successful hire, and low productivity. For positive experiences, for each hiring event, we calculate the share of each race that achieves tenure of at least one year in their job during our sample period, accounting for around 10% of hires. Long tenure suggests a successful hire and sufficiently good match between the worker

[11] Results are qualitatively and quantitatively similar comparing to race-specific turnover rates or aggregate turnover rates.

and the position. Given the potentially forward-looking nature of these measures, we exclude workers hired in the last 3 months (1 year) of our sample for negative (positive) experiences. Appendix C describes other features of workers associated with having a good or bad experience. The latter are never promoted in our sample, they have lower sales performance for subsets of workers for whom this information is available, and they are more likely to abandon their job without warning which is particularly costly for the firm, be terminated for unsatisfactory performance, and never get past the probation stage. In contrast, the former are more likely to leave for career advancement, studies, or personal reasons.

Summary statistics for performance measures are shown in the bottom panel of Table 4.1. Black hires have a higher probability of being terminated within 3 months, slightly lower probability of quitting within 3 months, and slightly lower probability of achieving at least one year of tenure. These differentials could indicate a lower average performance of black workers at the firm, but they are also endogenous to potential biases and discrimination. In any case, our theoretical framework's main predictions do not depend on true group productivity, but rather the difference between expected and realized performance for each race, both of which may be subjectively assessed.[12]

4.4 Results

Table 4.2 shows that a manager's previous experiences hiring worker groups have a clear impact on current hiring decisions. 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 and are better interpreted for a one standard deviation change than for extreme outcomes of 0 or 1. Estimated impacts in percentages are proportionally larger for black than white hiring given that they constitute a minority of workers, approximately 50% larger, indicating that hiring experiences

[12] If managers expect to fire black workers more often, then they may update less following such an event. In contrast, our empirical results indicate more updating, consistent with differential updating as discussed in our theoretical framework.

play a disproportionately large role in black hiring.

The first panel presents results for the cumulative measure of expected tenure by race, with more negative (positive) experiences resulting in a more negative (positive) measure of deviation from expected tenure and a higher (lower) hazard rate for a given tenure. The results show that a higher expected tenure for black workers based on previous experiences leads to statistically significantly more hiring of these workers, with a one standard deviation increase corresponding to an increase of 3-5% in the relative hiring of black workers. Columns 3-5 consider different samples which yield similar conclusions, focusing on the hiring of female workers, hiring by white managers only, and all hiring spells without restricting to new managers. The estimates for expected tenure with white workers are substantively smaller, not statistically significant, and tests reject the null that they are of equal magnitude but opposite sign to that of black workers.

The middle panel presents the results of negative experiences specifically. They indicate that managers decrease their hiring of black workers by around 6% for a one standard deviation increase in the fraction of previous black hires that were dismissed or quit within 3 months. Estimates across columns suggest that these impacts affect the hiring of both men and women, are not restricted to white managers, and are not simply driven by new managers. Across columns, estimates are statistically significant at the 5% level. Estimated impacts for experiences with white workers indicate a substantially smaller increase of approximately 2-3% in black hiring when accounting for the higher standard deviation of experience measures with black workers. Statistical tests reject the null that impacts with black and white workers are equal but of opposite sign for the baseline sample as well as that restricted to white managers, while differences are smaller for female workers and hires by more experienced managers. One possible explanation for larger responses following experiences with black hires is that proportionally more separations with black workers indicate dismissals rather than quits, but Appendix C shows that discrepancy in response is no smaller for quits. Moreover, there is a similar discrepancy in responses across groups following positive experiences, as we discuss next.

The bottom panel of Table 4.2 presents estimates of the impact of positive previous experiences. Managers statistically significantly increase their hiring of black workers by approximately 3-4% for a one standard-deviation increase in the fraction of previous black hires who reached at least one year of tenure in their position. Estimated impacts of experiences with white workers are smaller and statistically non-significant. Across panels and specifications, estimates suggest that cumulative previous experiences, particularly with black workers, have a substantial impact on subsequent hiring decisions, with some evidence that the discrepancy is smaller for female workers, little evidence that this is restricted to white managers, and evidence that these impacts do not only affect hiring decisions of new managers at the firm.

Negative experiences appear to have a larger impact than positive ones. There are several potential explanations for this, perhaps the most straightforward being differences in the performance measures.[13] A relatively quick dismissal or quit indicates a bad enough match that the employment relationship ended, while a worker remaining in their position only indicates sufficient performance and match quality to avoid separation. Other potential factors include turnover costs, risk aversion, and managers being more inclined to explain a negative outcome using external factors (the performance of a worker group) than their own performance. Still, the key takeaway is that both positive and negative experiences with groups, particularly black workers, impact hiring.

Appendix C presents additional results showing that both quits and fires have a negative impact on subsequent hiring, consistent with managers aiming to avoid both (Autor and Scarborough, 2008), that considering the performance of black hires relative to white hires or workers in the CBSA has limited impact on the results, and that there is a larger impact in areas with a larger black population, as predicted by the theoretical framework since belief updating is more likely to be pivotal in subsequent hiring decisions. Additional results also suggest that negative experiences with black workers carry over when a manager transfers stores, that negative experiences with black workers inherited by a new manager decreases hiring of the group by the manager, and that negative experiences which are less likely to be endogenous to a manager's behavior also decrease hiring.[14] Lastly, the cumulative average experiences

[13]We investigated other measures of positive experiences including promotions, salary and additional work hours. Promotions within the first months of employment are very rare, salaries are generally fixed, and additional hours primarily reflect demand fluctuations.

[14]These include dissatisfaction with pay, compensation or benefits, which are not controlled by

of other managers at the same store have little impact on a manager's own hiring decisions after accounting for the manager's own experiences. Managers learning from other managers could mitigate biases arising from a manager's own experiences. Yet, even in a setting where this type of information is relatively observable, informative given the similarity in employment contexts, and in which managers have incentives to cooperate given store-level bonuses, hiring decisions appear primarily driven by own experiences.

Accordingly, a manager's previous hiring experiences may affect their learning about groups by affecting their subsequent hiring. We turn to an analysis of the persistence of early experiences across multiple hiring events. The theoretical framework predicts that the impact of negative experiences should be more persistent if they lead to decreased hiring of a group, preserving potentially negatively-biased beliefs. Moreover, the impact may be more persistent for minority groups because of stronger belief updating, because it is easier for managers to avoid hiring from the group, and because they observe less information about the group from sources outside their own hiring. We estimate a similar regression model to equation (4.3), replacing the measure of cumulative average experience with measures of experience with the first group of black and white hires by a manager and the average of the first three groups.

The results are shown in Table 4.3. Estimates from the first panel indicate a statistically significant 1-2% decrease in black hiring for the current hiring event for a one standard deviation increase in the fraction of the first black hire(s) that were fired or quit within 3 months. Estimates from the second panel indicate a statistically significant decrease of 3-4% in subsequent black hiring for a one standard deviation increase in our measure of bad experience with the first three groups of black hires by a manager.[15] Column 4 includes hiring spells from more experienced managers which hired workers before the start of our sample, with the smaller coefficients suggesting that negative experiences at the beginning of a manager's career, rather than at the beginning of their current hiring spell, are particularly impactful. In contrast, impacts of early negative experiences with white worker(s) are smaller and

the department manager, as well as worker integrity, illegal or unethical behavior, or violation of rules and policies.

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

statistically non-significant. Panels 4 and 5 also show smaller and generally statistically non-significant impacts for early positive experiences with both black and white workers, highlighting a different persistence pattern from early negative experiences with black workers in particular.

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

Lastly, we investigate how managers' most recent experiences affect their hiring, since they could be particularly salient due to recency bias or evolving hiring contexts and applicants pools. The third panel of Figure 4.3 indicates a statistically significant decrease of around 5% in current hiring of black workers for a one standard deviation increase in the fraction of the latest black worker(s) hired by the manager that quit or were fired within 3 months. Estimates of the impact of the latest experience with white workers are approximately 30% smaller, but also statistically significant. Panel 6 presents results for positive experiences corresponding to a 2% increase (1.5% decrease) in black hiring for a one standard deviation increase in our positive experience measure for the latest black (white) hire(s). These results indicate that recent experiences with both worker groups are salient, including for more experienced managers as shown in column (4).

4.4.1 Interpretation and Alternative Explanations

Overall, our results broadly support the predictions of our theoretical framework. First, cumulative previous experiences of managers with worker groups impact current hiring of these groups in the predicted direction. Second, early negative experiences, particularly with black workers, have more persistent impacts than positive ones. Third, managers react more following hiring experiences with black workers. As such, minority workers are disproportionately impacted, because of differential updating by managers combined with the relative persistence of negative experiences on subsequent hiring for these workers.

Managers learning about the performance of worker groups from their experience is one natural explanation for these findings. Persistence in hiring patterns across groups could also be driven by supply-side adjustments, such as workers selectively applying for positions with managers based on their hiring record. These provide a poor alternative explanation: workers apply for a job at the store or area level, do not know their manager until the interview, and are unlikely to observe information about the manager's hiring record until they are employed.[16] Employment through referrals with racially-homogeneous networks also provides a poor alternative explanation. First, we document larger hiring responses following experiences with black workers although evidence suggests that proportionally more white workers are hired through referrals (Kirnan et al., 1989; Taber and Hendricks, 2003). Second, for referrals to be effective, managers should hire workers referred by high-performing workers, inconsistent with decreases in black hiring following negative experiences being driven by foregone referrals. Third, positive experiences as measured do not reflect a direct change in worker group composition, but do affect subsequent hiring. Fourth, our results along with institutional details suggest that workers are not assigned to department managers based on their performance with groups, especially since we exclude transfers.[17]

Selection on unobservables and taste-based discrimination (including by co-workers), such that negative (positive) experiences reflect a bad (good) working climate for minority workers which translates to less (more) hiring, are also poor alternative explanations. To the extent that work climate and taste-based discrimination are relatively time invariant, they are captured by store and manager fixed effects. Further, it seems unlikely that negative experiences with white workers reflect prejudice or a bad working climate for white workers in predominantly white departments with white managers. Yet, we find that experiences with white workers do impact black hiring, and vice versa. Lastly, such alternatives provide little rationale for the persistence of negative experiences in particular.

[16] Managers may treat workers differently during the hiring process in a way that makes black workers less likely to receive an offer or accept it, but that is broadly consistent with our mechanism: negative experiences create biases which shape views and behavior in a self-sustaining way.

[17] Similarly, we find little evidence that a higher fraction of black workers in a manager's current team leads to more black hires in the future regardless of experiences and after controlling for store and manager fixed effects, conflicting with worker-group complementarities being an alternative explanation. These would also not explain why experiences with black workers have larger, more persistent impacts on hiring.

Managers learning about their own ability rather than worker groups is also unlikely to be driving our results, and whether some managers are indeed better or worse at hiring certain groups is inconsequential. Appendix C suggests that inheriting worse black workers decreases subsequent black hiring, though precision is low, and experiences with black workers who were fired or quit for reasons unlikely to be related to the manager's performance also decrease subsequent black hiring. There is also extensive evidence in psychology and economics that managers, employers, and CEOs are routinely overconfident and attribute negative outcomes to external sources (Moore et al., 2015; Guenzel and Malmendier, 2020).

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

4.5 Conclusion

We study the determinants of individual manager heterogeneity in the hiring of racial worker groups using the employment records of a large US retailer and studying repeated hiring decisions of managers. We find that the hiring context, such as the location of a store, is an important determinant of the hiring of black workers across managers, but so are manager fixed effects and previous hiring experiences of the manager with these workers. Consistent with a theoretical framework which combines statistical discrimination with learning about worker groups, we find that 1) negative (positive) experiences with black and white workers decrease (increase) subsequent relative hiring of these groups, 2) the impact of experiences with black workers on subsequent hiring is disproportionately large, and 3) negative experiences with black workers lead to particularly persistent decreases in subsequent hiring of the group.

Our results have a number of implications for the literature, including for the study of discrimination and the organization of the firm. We contribute to a growing body of work on managers, particularly in decentralized organizations, having discretion in hiring which leaves room for individual biases (Hoffman et al., 2018; Berson et al., 2019; Benson et al., 2020; Bergman et al., 2020). Since inefficient hiring is costly for the firm, this suggests scope for the design of organizational policies. Some options include pre-employment testing and the use of hiring algorithms, but the firm we study already has pre-employment testing and the patterns we document can arise even if managers are profit-maximizers and Bayesian over their own experiences. More targeted information aggregation and sharing between managers may be necessary.

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

Figure 4.1: Manager Heterogeneity in the Hiring of Black and White Workers

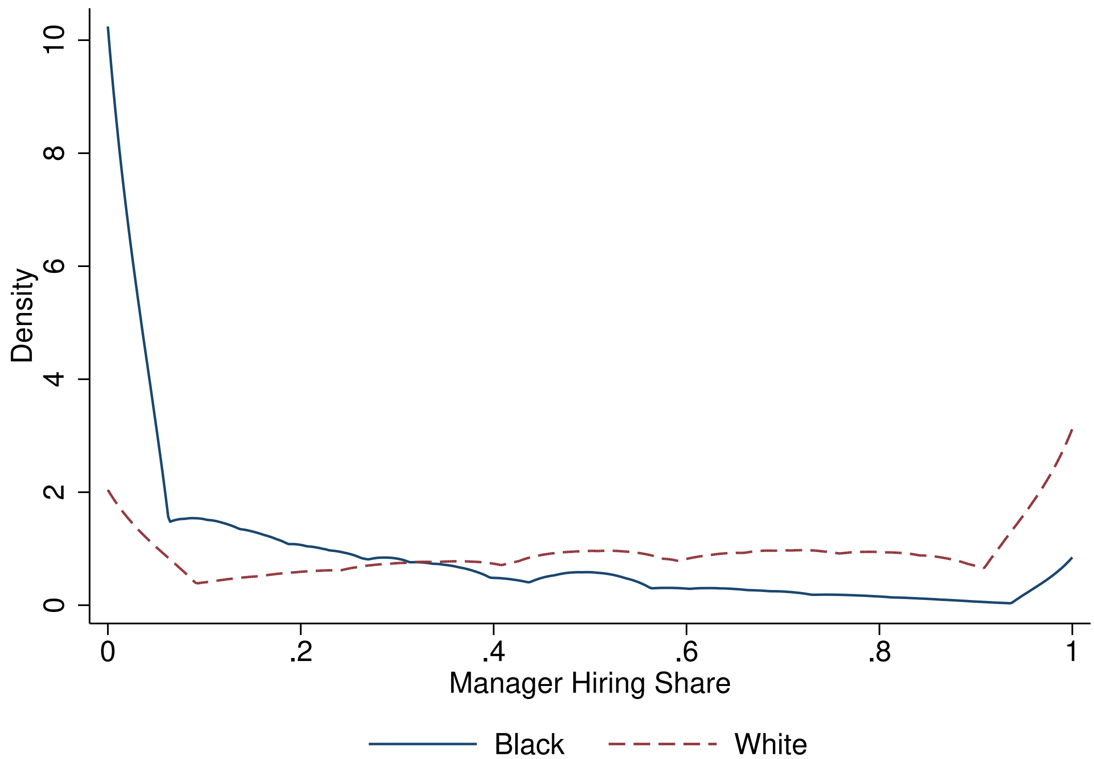

Figure 4.2: Manager Predicted Shares of Black Hiring Based on Hiring Context, Market Factors, and Manager Fixed Effects

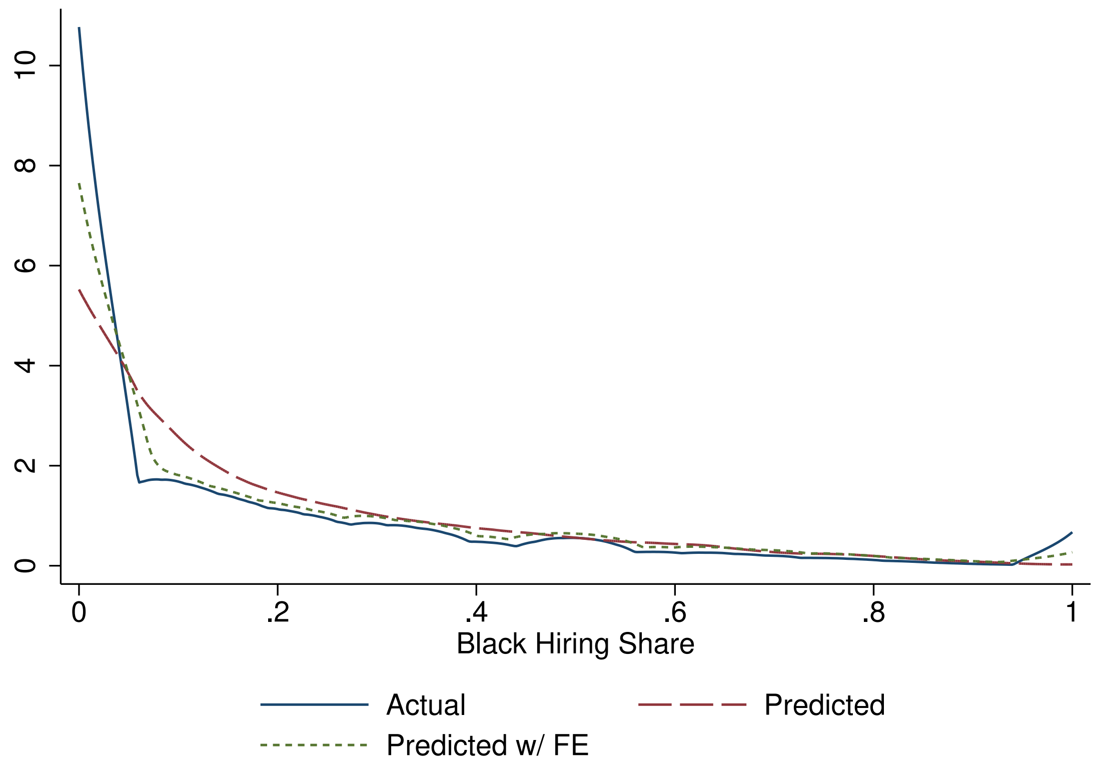

Predicted shares are obtained by averaging predicted values for each manager from an individual hire level linear probability model regression including whether the worker was hired for a part-time or full-time job, the manager's previous number of hires at the time that the current worker is hired, yearly state unemployment rate and fraction with at least some college education, and month and year, store, job title, and individual manager fixed effects.

Figure 4.3: Distribution of Manager Fixed Effects for Black Hiring

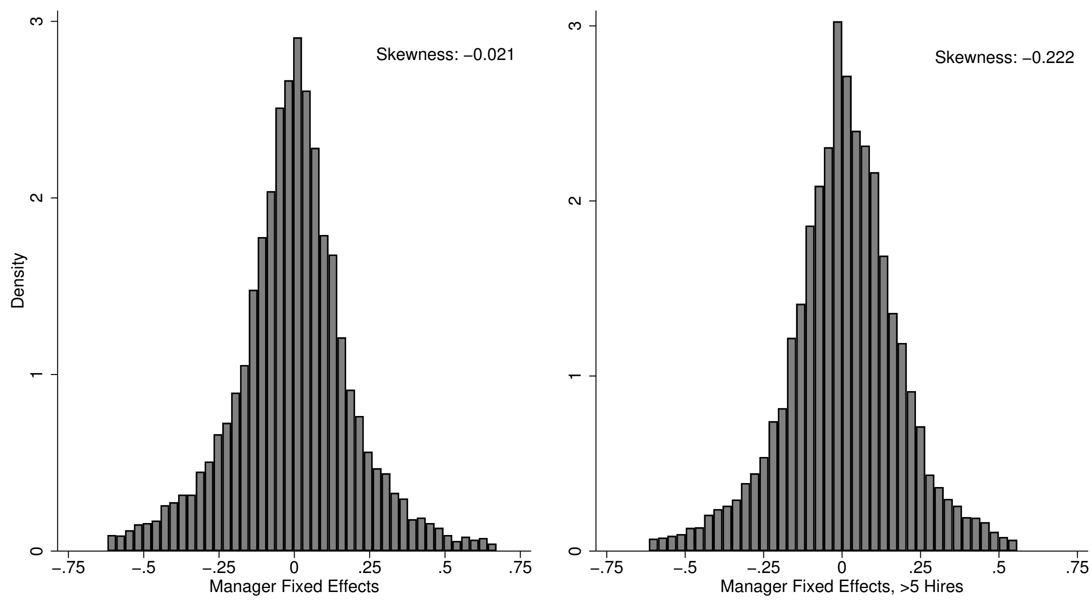

See Figure 4.2 for specification details. Fixed effects are estimated for the largest connected sample of stores and managers following Abowd et al. (1999) and adjusted using empirical Bayes shrinkage.

Table 4.1: Summary Statistics and Performance Measures

	Workers	Managers
Age	31.01	41.82
	(14.14)	(11.39)
Female	0.56	0.37
	(0.50)	(0.48)
White	0.55	0.73
	(0.50)	(0.44)
Black	0.22	0.11
	(0.41)	(0.31)
Hispanic	0.16	0.12
	(0.37)	(0.32)
Other	0.07	0.05
	(0.26)	(0.21)
Tenure	33.73	122.75
	(71.68)	(125.17)
Full Time	0.18	0.99
	(0.37)	(0.09)
# Hires		30.16
		(65.75)
N. Obs	1,067,682	27,470

Performance Measures	Black	White
Fired within 3 months	0.768	0.044
	(0.266)	(0.205)
Quit within 3 months	0.188	0.198
	(0.391)	(0.399)
Tenure above 1 year	0.089	0.108
	(0.284)	(0.311)

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

Table 4.2: Cumulative Impact of Previous Experiences with Black and White Workers on Current Hiring of Black Workers

			Female Workers	White Managers	All Spells
Expected Tenure					
Black	0.048	0.050	0.041	0.044	0.060
	(0.011)	(0.011)	(0.016)	(0.016)	(0.008)
White		-0.009	0.023	-0.001	-0.012
		(0.011)	(0.017)	(0.013)	(0.007)
Outcome Mean	0.380	0.380	0.439	0.326	0.365
Observations	35,937	35,937	22,312	22,510	72,424
P-Value: Black = -1 * White		0.023	0.004	0.053	0.000
Fraction Quit or Fired					
Black	-0.073	-0.072	-0.057	-0.059	-0.067
	(0.016)	(0.017)	(0.022)	(0.019)	(0.011)
White		0.044	0.031	0.017	0.067
		(0.022)	(0.027)	(0.025)	(0.016)
Outcome Mean	0.367	0.348	0.403	0.307	0.337
Observations	33,971	31,911	19,546	23,741	66,692
P-Value: Black = -1 * White		0.077	0.310	0.060	0.169
Fraction Long Tenure					
Black	0.062	0.058	0.072	0.043	0.062
	(0.023)	(0.024)	(0.031)	(0.028)	(0.016)
White		-0.001	-0.047	0.012	-0.015
		(0.027)	(0.038)	(0.030)	(0.018)
Outcome Mean	0.366	0.347	0.402	0.306	0.336
Observations	28,456	26,655	16,198	19,869	56,911
P-Value: Black = -1 * White		0.090	0.522	0.155	0.025

Expected Tenure corresponds to the cumulative average deviation from expected tenure at the firm for workers hired by the manager. Regressions include the fraction of full-time and female hires, average age of hires, total number of workers hired in the event, number of previous hiring events, time since the last hiring event, yearly unemployment and college attainment rates in the state, month and year, store, and manager fixed effects. Clustered standard errors at the manager level are presented in parentheses. Column 3 restricts the sample to female hires, column 4 restricts the sample to white managers, column 5 includes managers who were hiring at the firm before the start of our sample.

Table 4.3: Persistence of Previous Experiences with Black and White Workers on Current Hiring of Black Workers

		Female Workers	White Managers	All Spells
Fraction Quit or Fired				
First Event				
Black	-0.019	-0.024	-0.016	-0.007
	(0.007)	(0.009)	(0.008)	(0.004)
White	-0.009	-0.015	-0.013	0.000
	(0.008)	(0.011)	(0.010)	(0.005)
Outcome Mean	0.367	0.418	0.328	0.354
Observations	35,613	22,013	26,404	72,880
P-Value: Black = -1 * White	0.006	0.005	0.023	0.198
Average of Events 1-3				
Black	-0.042	-0.060	-0.033	-0.024
	(0.014)	(0.020)	(0.018)	(0.009)
White	-0.013	-0.046	-0.023	-0.008
	(0.017)	(0.022)	(0.023)	(0.011)
Outcome Mean	0.402	0.449	0.337	0.385
Observations	27,829	17,417	22,177	59,762
P-Value: Black = -1 * White	0.005	0.000	0.030	0.010
Previous Event				
Black	-0.041	-0.043	-0.039	-0.043
	(0.006)	(0.008)	(0.007)	(0.004)
White	0.035	0.029	0.032	0.036
	(0.006)	(0.007)	(0.007)	(0.004)
Outcome Mean	0.348	0.403	0.307	0.337
Observations	31,911	19,546	23,741	66,692
P-Value: Black = -1 * White	0.174	0.1297	0.168	0.020
Fraction Long Tenure				
First Event				
Black	-0.010	0.013	0.009	-0.012
	(0.011)	(0.014)	(0.014)	(0.006)
White	-0.003	0.005	0.011	0.002
	(0.011)	(0.014)	(0.013)	(0.006)
Outcome Mean	0.366	0.4170849	0.328	0.353
Observations	29,869	18,324	22,220	62,487
P-Value: Black = -1 * White	0.375	0.3559	0.268	0.271
Average of Events 1-3				
Black	0.020	0.057	0.013	0.006
	(0.018)	(0.030)	(0.026)	(0.012)
White	0.018	-0.013	0.012	0.002
	(0.020)	(0.029)	(0.028)	(0.012)
Outcome Mean	0.401	0.446	0.364	0.385
Observations	23,527	14,736	17,268	51,696
P-Value: Black = -1 * White	0.146	0.298	0.489	0.645
Previous Event				
Black	0.022	0.026	0.021	0.019
	(0.009)	(0.012)	(0.010)	(0.006)
White	-0.016	-0.021	-0.013	-0.006
	(0.008)	(0.011)	(0.009)	(0.005)
Outcome Mean	0.347	0.402	0.306	0.336
Observations	27,249	16,198	19,869	58,026
P-Value: Black = -1 * White	0.658	0.7575	0.558	0.080

Clustered standard errors at the manager level are presented in parentheses. First Event refers to the first group of workers from a given race hired by the manager, Average of Events 1-3 refers to the average outcome of the first three groups (excluding managers who hired less than three groups of workers from either racial group). Latest Event refers to the latest group of hires from a given race. See Table 4.2 for additional details.

CHAPTER V

Conclusion

This dissertation presents a new theory of discrimination in the labor market based on employers developing biased beliefs about the productivity of worker groups.

Chapter II presents a statistical discrimination model in which employers are initially uncertain about the productivity of worker groups, perceive group membership as a potentially relevant indicator of individual productivity, and endogenously learn about group productivity through their own hiring. In a dynamic setting, employers have noisier initial information on one group's productivity relative to another and trade off learning about that group against current-period profit maximization. An employer's hiring history shapes their future hiring and learning. Positive experiences create positive biases, which endogenously correct themselves through more hiring and learning. Negative experiences, however, create negative biases which decrease hiring and learning. Differential learning across employers results in a negatively-skewed distribution of beliefs about the group's productivity which can cause the wage of the group about whose productivity employers have noisier initial information to fall and remain below their expected productivity in the long run. The model generates steady state predictions analogous to taste-based discrimination, with endogenous beliefs replacing preferences, providing a new way to understand prejudice as the result of life experiences shaping beliefs in distortionary ways.

In Chapter III, I create a controlled environment to study how biased beliefs arise through endogenous employer learning about groups. Workers perform the real-effort task of solving character puzzles on the computer, which corresponds to their productivity in the context of the experiment. Employers then repeatedly hire workers, choosing between one of the two groups each period and observing their hire's productivity. They are incentivized to hire the most productive workers available, requiring them to identify which group is more productive, if any. I study how negative biases about group productivity arise from an employer's hiring history. I give employers better initial information on the productivity of one group and focus on hiring and learning about the other group. Consistent with the theory, the results show that negative hiring experiences, captured through the hiring of relatively low productivity workers, lead to persistent negatively-biased beliefs about the group's productivity, specifically by decreasing subsequent hiring of the group and therefore learning. In contrast, positive experiences increase subsequent hiring and learning, mitigating positive biases. Across employers, differential hiring and learning generate a persistent negatively-skewed distribution of beliefs about the group's productivity.

In Chapter IV, I use administrative data from the U.S. operations of a large national retailer to examine how managers' past experience hiring black and white workers affects the race of their subsequent hires. We document three specific hiring patterns that are particularly consistent with endogenous employer learning about worker group productivity. First, positive past experiences with white or black workers increase a manager's propensity to hire from that group, whereas negative experiences decrease their propensity. Second, learning is asymmetric by race. Although past experience affects subsequent hiring for both black and white workers, the effect is particularly pronounced for black workers. Third, early negative hiring experiences lead to a substantial persistent decline in black hiring, while the effects of early experiences with white workers or early positive experiences with black workers are comparatively small and short-lived.

APPENDICES

APPENDIX A

Appendix for Chapter II

A.1 Proofs of Propositions 1-3

Proposition 1

By market clearing, the marginal employer is indifferent between hiring from either group, implying $\lambda_{mt} = w_{Bt}(\Psi_t) - w_A$. Define $\lambda_{mt} = \lambda_t^c$. Given current beliefs and wages, profit maximization implies that employers with $\lambda_{jt} > \lambda_t^c$ strictly prefer to hire from group B while those with $\lambda_{jt} < \lambda_t^c$ strictly prefer to hire from group A. Thus, λ_t^c represents the cutoff relative WTP for a group B worker in period t.

Proposition 2

Given the prior $\mu_B \sim N(\mu_0, 1/\tau_0)$ and i.i.d hiring signals $x_1, ..., x_K$ drawn from $X \sim N(\mu, 1/\tau)$, the Bayesian Central Limit Theorem implies under standard regulatory conditions that the posterior belief distribution converges in distribution to μ as $K \to \infty$. The posterior distribution of beliefs for employers who remain above the hiring cutoff in the long run converges in distribution to μ. For almost all of these these employers, this implies that the value of learning converges to 0 such that $\lambda_{jt} \to 0$ as $K \to \infty$. Market clearing requires that a subset of employers hire from group B, for almost all of whom $\lambda_{jt} \to 0$ as $K \to \infty$ and $\lambda_{jt} \geq w_{Bt}(\Psi_t) - w_A$. Thus, $w_A \geq w_{Bt}(\Psi_t)$ asymptotically.

Market clearing also requires that a subset of employers hire from group A asymptotically, implying $\lambda_{jt} \leq \lambda_t^C$ for those employers. Define

$$\Delta V_{jt} = E_t[V(\psi'_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))] - E_t[V(\psi_{S_{jt+1}}, w_{Bt+1}(\Psi_{t+1}))] \text{ and}$$
$$\Delta f_{jt} = \mu - E[\mu_B|S_{jt}].$$

Employer j hires from group A only if $\beta\Delta V_{jt} - \Delta f_{jt} \leq w_{Bt}(\Psi_t) - w_A$. Since the value of information ΔV_{jt} is weakly positive, then $\Delta f_{jt} > 0$ for this group. $\Delta f_{jt} > 0$ implies that $E[\mu_B|S_{jt}] < \mu$. Employers who hire from group A asymptotically must have negatively-biased beliefs. Let F_B denote the fraction of group B workers. Asymptotically, since unbiased employers hire from B and biased employers from A, the fraction of biased employers is equal to $1 - F_B$ by market clearing.

Proposition 3

First, I show that w_{Bt} is strictly decreasing in t. Define E_{Bt} as the subset of employers who hire from group B in a given period t, with the fraction of employers in E_{Bt} equaling F_B. By definition, $\lambda_{jt} \geq w_{Bt} - w_A$ for these employers. Given a continuum of employers, some employers arbitrarily close to the cutoff observe a low signal, such that there exists $e_{Bt+1} \subset E_{Bt}$ with $\lambda_{jt+1} < w_{Bt} - w_A \leq \lambda_{jt}$.[1] Suppose $w_{Bt+1} \geq w_B$, then $E_{Bt+1} \subset E_{Bt}$ and the labor market doesn't clear. Thus, w_{Bt+1} must be smaller than w_{Bt} for all t.

Second, I show that $w_{Bt} \to c \in \mathbb{R}$ as $t \to \infty$. Since w_{Bt} is strictly decreasing in t, this is equivalent to establishing that w_{Bt} cannot fall below an arbitrarily low limit $\underline{\mathrm{w}}$. In any period, even asymptotically, employers below the hiring cutoff have observed a finite number of signals (if any). Then, they have a strictly positive value of learning about group B and posterior mean beliefs strictly above negative infinity. Denote $\lambda_{\underline{j}} = \underline{\mathrm{w}} > -\infty$ where $\lambda_{\underline{j}}$ is the supremum relative WTP for a group B worker for

[1] This does not rely on unbounded signals. The continuum assumption ensures that a mass of employers is arbitrarily close to the cutoff each period.

employers below the cutoff as $t \to \infty$. Then, $w_{Bt} \geq \underline{\mathrm{w}}$ for any t. Since w_{Bt} is strictly decreasing in t but bounded below, it must converge to a constant as $t \to \infty$.

Third, I show that $c < w_A$. For any $\varepsilon > 0$, there exists a t large enough such that fraction $F_B - \varepsilon$ of employers currently hiring from Group B have value of learning smaller than ε and will hire from Group B in the limit.[2] There also exists $t' > t$ arbitrarily large such that beliefs of employers hiring from Group B at t' are almost entirely driven by signals observed between t and t'. More precisely, $\mu_B|S_{t'j}$ follows approximately the same distribution as $\mu_B|\{S_{t'j} \setminus S_{jt}\}$ with the same parameters. Given that $E[\mu_B|\{S_{t'j} \setminus S_{jt}\}]$ converges to μ for almost all employers who hire from group B, some employers who hire from group B at t' have posterior mean beliefs below μ[3] and a value of learning smaller than ε, such that their relative WTP for a group B worker λ_{jt} is below 0. By market clearing, the relative WTP of the marginal employer is no greater than the infimum relative WTP of employers hiring from group B, implying that $\lambda_{mt} = w_{Bt} - w_A < 0$ and thus that $w_{Bt} < w_A$ for $t > t'$. Since w_{Bt} is strictly decreasing in t, then $c < w_A$.

A.2 Additional Model Implications

Signals of Individual Productivity

Consider the case in which employers observe a noisy signal s_i of individual worker productivity x_i at the hiring stage and do not rely solely on group membership g to predict productivity. This signal is exogenous, rather than the result of an investment choice, and can be thought of as a score on a pre-employment test. Employers observe

$$s_i = x_i + \varepsilon_i$$

[2]This is because the value of learning and the probability that an employer currently hiring from group B falls below the cutoff next period go to 0 asymptotically.

[3]The probability that the posterior beliefs of employers all converge in distribution to μ from above is 0 given a large number of employers and signals.

where $\varepsilon_i \sim N(0, 1/\tau_\varepsilon)$ is i.i.d. random noise. They estimate productivity according to the following rule

$$E[x_i|s_i, S_{jt}] = \gamma s_i + (1-\gamma)E[\mu_g|S_{jt}]$$

where $\gamma_{gjt} = \frac{1/\tau + Var[\mu_g|S_{jt}]}{1/\tau + Var[\mu_g|S_{jt}] + 1/\tau_\varepsilon}$ is a measure of the signal's precision. Negatively-biased beliefs about the mean productivity of group B arise as in the baseline model. Since employers above the hiring cutoff are willing to pay more for a group B worker conditional on a given signal value, workers and employers sort such that hiring and learning dynamics are also unchanged. Workers can be indexed by their signal value, with the same learning problem arising for each worker "type" and a market-clearing wage for each type-group pair.

Discrimination may still vary by occupation, skill, and education depending on the variance in productivity and productivity signals. These variances determine the extent to which employers rely on group membership to predict productivity, and therefore the importance of the learning problem. Discrimination empirically appears smaller for high-skill workers, at least in the case of race (Lang and Lehmann, 2012). Differences in the information available at the time of hiring, variance in productivity, or the speed with which the market learns individual worker productivity, could all help explain this empirical regularity (Arcidiacono et al., 2010; Lindqvist and Vestman, 2011).

Endogenous Worker Investments

When groups are ex-ante equally productive, statistical discrimination models usually generate outcome disparities by showing that workers from group B may face different incentives to invest in human capital, for example due to employers perceiving their signals of productivity as noisier (Lundberg and Startz, 1983) or because they hold negative stereotypes against them (Coate and Loury, 1993). Statistical discrimination therefore arises when group B becomes less productive due to lower investment.

In my model, even if employers have biased beliefs on average, workers and employers sort such that group B is hired by employers above the cutoff who have approximately unbiased average beliefs with experience. Accordingly, group B doesn't necessarily have incentives to invest differentially in human capital due to biased beliefs of employers. Nevertheless, group B may expect a different return for the same investment if relative wages across investment levels vary due to the nature of individual signals of productivity. Group B workers may be incentivized to sort into areas or occupations where the information asymmetry problem faced by employers is lesser, providing a rationale for group specialization. Similarly, if group B earns lower returns from the labor market overall, they may have incentives to invest less in human capital which could exacerbate discrimination.

Firm Size

Employers who hire more workers have a higher value of learning and may learn more quickly. Negative biases may be less likely to arise and persist, and these employers may hire a higher fraction of group B workers, consistent with evidence reported in Holzer (1998) and Miller (2017) for black workers. These implications presumably relate to large establishments with centralized, professional human resources (HR) departments rather than large firms with decentralized hiring across smaller establishments.[4]

Implications for market-level discrimination are limited if each establishment hires a negligible fraction of workers and there is size heterogeneity above the hiring cutoff. Unless all of group B is hired by large establishments with centralized hiring, these establishments are not marginal, by definition, and the wage is determined by smaller establishments. Casual empiricism certainly suggests that some small firms and large firms with decentralized hiring hire workers from groups typically of interest in the discrimination literature. A back of the envelope calculation suggests that around 17% of black workers were employed at firms with less than 25 workers in 1998, and this proportion is substantially larger for establishments under 25 workers.[5]

[4] Evidence from hiring at decentralized firms suggests that individual managers play an important role in the racial composition of hires (Giuliano et al., 2009; Giuliano and Ransom, 2013; Benson, Board and Meyer-ter-Vehn, 2019).

[5] This is based on Headd (2000) which provides the proportion of black workers across firm size in 1998 combined with statistics from the Census Bureau on the total number of workers employed

Search Frictions

A formal search model is beyond the scope of this paper, but previous work suggests that search frictions may have important implications. With random search, the intuition behind the endogenous learning mechanism is vastly unchanged. Employers who hire from group B and have negative (positive) experiences are less (more) likely to select a worker from the group again in the future. Positive biases are learned away more quickly than negative ones, so the distribution of beliefs is negatively skewed. The wage gap would be determined by the average rather than the marginal employer, as highlighted by Black (1995) in reference to the Becker model. Accordingly, wage gaps along the equilibrium path may be larger in a search framework. Search frictions could also mitigate the stark prediction that employers below the hiring cutoff in the long run never hire from group B again.

Minority Employers

The role of group B employers depends on whether they share the beliefs of group A employers or face the opposite learning problem (know group B productivity and learn about group A). In the first case, the distinction between employer groups is irrelevant for my purpose. In the second case, these employers constitute a fraction of the market who may not develop biases about group B, encouraging segregation and mitigating wage gaps. Other factors may disadvantage group B employers: uncertainty about the majority group or discrimination in promotion and the capital market (Farrell et al., 2020). Empirical evidence for both race and gender suggests that the proportion of managers is low compared to that of workers (Giuliano et al., 2009; Blau and Kahn, 2017).

A.3 Simulations and Comparative Dynamics

Given a prior distribution of beliefs, the initial market-clearing wage when employers maximize their expected profits is found. Beliefs are updated such that those above

at firms below 25 workers and the total number of black workers for the same year.

the cutoff receive a signal of productivity from group B and others retain their beliefs. Given this new distribution of beliefs, a new market-clearing wage is found, and the process is repeated. The dynamic optimization problem is solved for a discretized state space which gives the value of learning for combinations of beliefs and wages through interpolation. Worker productivity is distributed $N(0, 2)$ and prior beliefs are distributed $N(0, 1)$. The group A wage w_A is normalized to 0 and the discount factor β is set to 0.9.

The expected size of the wage gap is influenced by the exogenous parameters of the model as in Figure A.1. Namely, as in the Becker model, a higher fraction of group B workers is predicted to lead to a lower wage for group B. A lower mean group productivity also leads to a lower wage. If employers have negatively-biased priors about group B productivity, then their wage will be lower initially and reach a similar level in the long run. A higher prior precision or lower variance in productivity increases the wage of group B. Assuming homogeneous rather than unbiased priors has little impact on the wage (slightly higher), while introducing stereotype bias through employers overestimating the precision of their signals (or equivalently underestimating the variance in group B's productivity) decreases the wage. With entry and exit of employers, when new employers hold unbiased priors, a lower exit rate differential for employers who hire from group A leads to a lower wage for group B, as does a lower aggregate exit rate.

In Figure A.2, I conduct simulations with entry and exit of employers in which employers below the cutoff are 100% more likely to exit in any given period and new employers enter the market with biased priors. Namely, mean prior beliefs equal the average belief of employers already in the market. A wage gap is sustained and larger than when employers have unbiased priors as in Figure A.2, even with a higher differential exit rate.

Figure A.1: Wage Gap and Model Parameters

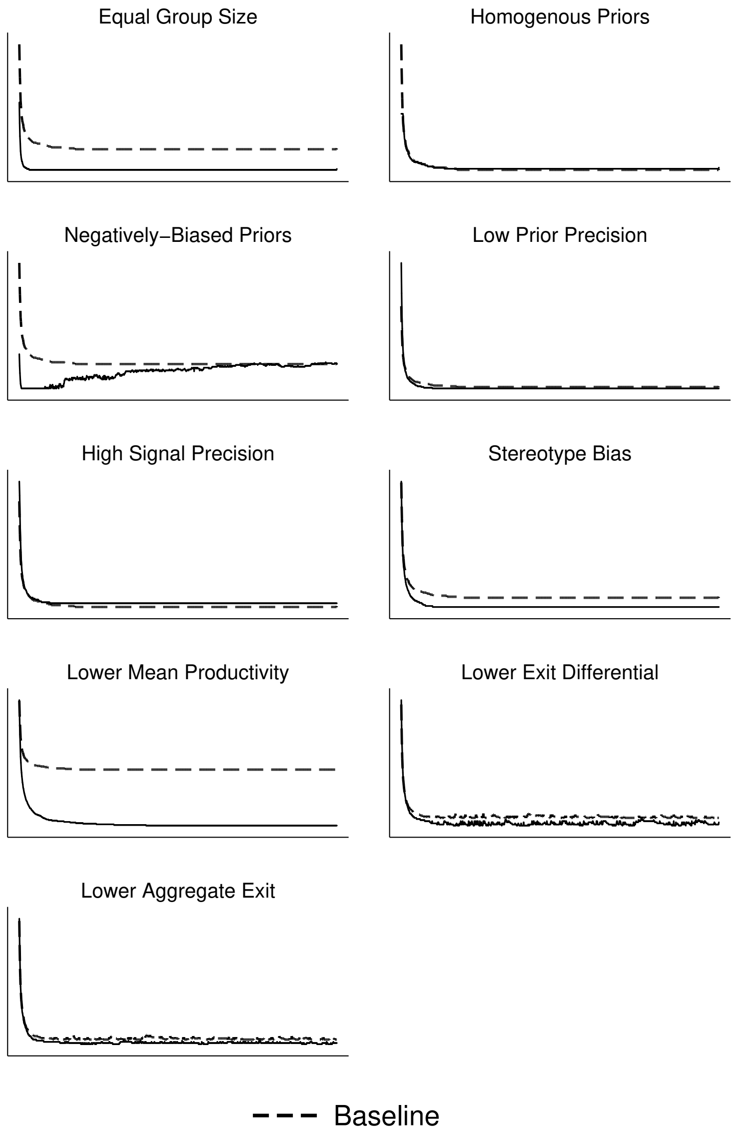

Equal Group Size refers to group B being of equal size to group A (50% of workers). Homogenous Priors refers to each employer holding prior $\mu_0 = 0$. Negatively-Biased Priors refers to employers having mean prior beliefs below the true value (-1 vs 0). Low Prior Precision corresponds to a case with prior variance equal to 2. High Signal Precision corresponds to a case with variance in worker productivity equal to 1. Stereotype bias corresponds to a case where employers incorrectly believe group B worker productivity to be 2 when it is 4. Lower Mean Productivity corresponds to a case where mean group B productivity is lower than that of group A (-1 vs 0). Lower Exit Differential refers to a case where employers who hire from group A are 10% more likely to exit the market each period. Lower Aggregate Exit refers to a case where the overall exit rate is 1% each period. See Figure A.1 for other parameter choices.

Figure A.2: Model Simulation with Market Entry and Exit, 100% Exit Differential, Biased Priors

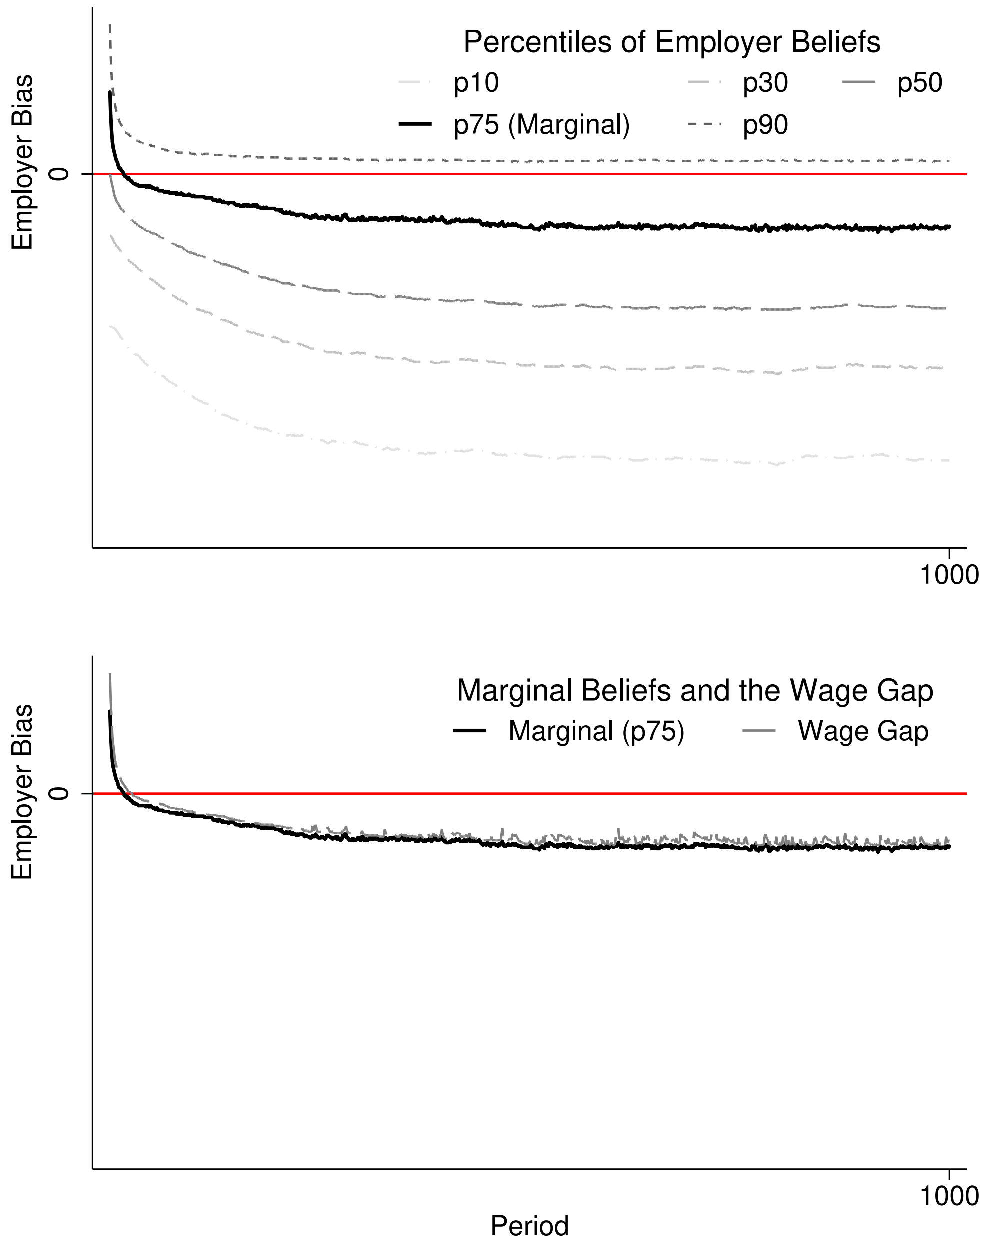

The aggregate exit rate corresponds to 2% each period, with a 100% higher exit rate for employers below the hiring cutoff for group B. New entrants have mean beliefs equal to the mean of employers already in the market. See Figure A.1 for other parameter choices.

A.4 General Productivity Distribution

Let worker productivity be drawn from $X|\mu_B \sim G(x)$, a one-parameter family of distributions characterized by their mean, with full support on an interval of real numbers $\mathbb{X}$, bounded variance, and density function $g(x)$. The parameter of interest is the expected productivity of group B, $\mu_B = E_G[x]$. Employers have a common prior distribution about group B's mean productivity $h(\mu_B)$. Each hire provides an i.i.d. private signal x about worker productivity and S_{jt} is the collection of all signals observed by time t. Under strictly monotone and continuous Bayesian updating on the mean, the distribution of posterior beliefs conditional on S_{jt} corresponds to

$$z(\mu_B|S_{jt}) = \frac{\Pi_{k\epsilon S_{jt}} g_{x_k}(x_k)h(\mu_B)}{\int \Pi_{k\epsilon S_{jt}} g_{x_k}(x_k)h(\mu_B)d\mu_B}.$$

The hiring decision hinges on the expected productivity of Group B, which decreases with lower signals about the group's productivity. As such, hiring decisions, market clearing conditions and wage setting are unchanged, along with Proposition 1. Proposition 2 follows under regularity conditions on $G(\cdot)$ and $h(\cdot)$. Proposition 3 follows from assumptions made on $G(\cdot)$ as well as Propositions 1-2.

APPENDIX B

Appendix for Chapter III

B.1 Example Puzzle

Figure B.1: Example Puzzle

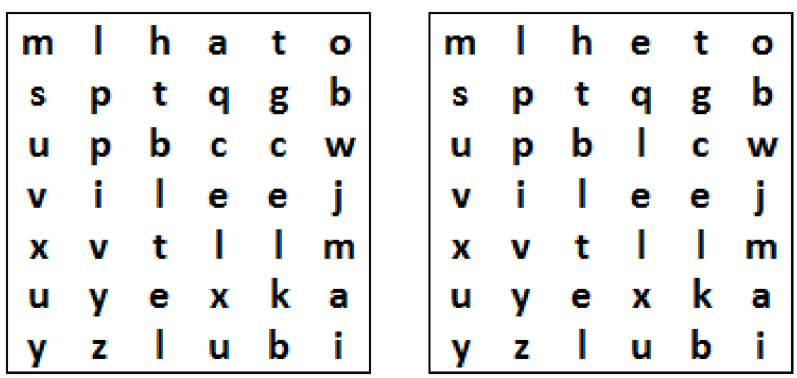

The square with characters on the right differs from the square on the left in two letters. Workers had to identify those letters to solve the puzzle.

B.2 Experimental Instructions

B.2.1 Workers

Page 1. There are two roles in this study: workers and employers. You will play the role of the worker. Your task is to solve as many character puzzles as possible within a 4:00 minute period. You will be able to solve a practice puzzle on the next page to familiarize yourself with the task. For each puzzle that you solve during the four minutes, you will receive 250 credits. For example, if you solve 10 puzzles you will receive 250 x 10 = 2500 credits.

Your performance on the task will be recorded and added to a pool of workforce available to employers. Employers will be tasked with selecting the best workers from the pool. When solving puzzles, you can only continue to the next page once you enter the correct answer.

Before payment, your performance may be evaluated. If the study was not completed with reasonable effort, such as if no puzzles were solved, no payment will be made.

Page 2. On this page, you have the opportunity to solve an example puzzle to familiarize yourself with your task. The square with characters on the right differs from the square on the left in two letters. You have to find those letters and enter them in the submission box to solve the puzzle. For your submission to be valid, you must enter the two letters from the square on the RIGHT without spaces in the order in which they appear going row by row and then from left to right.

m	l	h	a	t	o
s	p	t	q	g	b
u	p	b	c	c	w
v	i	l	e	e	j
x	v	t	l	l	m
u	y	e	x	k	a
y	z	l	u	b	i

m	l	h	e	t	o
s	p	t	q	g	b
u	p	b	l	c	w
v	i	l	e	e	j
x	v	t	l	l	m
u	y	e	x	k	a
y	z	l	u	b	i

In this case, the fourth letter of the first row is "e" in the box on the right but "a" in the box on the left and the fourth letter of the third row is "l" in the box on the right but "c" in the box on the left. Accordingly, the correct answer is "el". Note that entering “le” will not count as a correct answer. You can submit answers by pressing the enter key or clicking on the ”next” button.

Page 3. On the next page, you are asked to solve a timed practice puzzle. Your time will be recorded and may be visible to employers who will later be tasked with selecting workers.

Only go to the next page when you are ready. The practice puzzle will start immediately.

Page 4. Solve the puzzle below:

z	g	p	r	d	f
j	x	d	a	f	l
l	m	q	d	i	v
i	s	b	t	s	p
n	p	m	w	d	q
l	x	b	b	m	t
p	k	n	v	u	z

z	g	p	r	d	f
j	x	d	a	f	l
l	m	q	d	i	v
i	s	b	d	s	p
n	p	m	w	d	q
l	x	b	s	m	t
p	k	n	v	u	z

Reminder: the square with characters on the right differs from the square on the left in two letters. You have to find those letters and enter them in the submission box to solve the puzzle. For your submission to be valid, you must enter the two letters from the square on the RIGHT without spaces in the order in which they appear going row by row and then from left to right.

Page 5. Only go to the next page when you are ready. Your task will begin immediately and you will have 4 minutes to solve as many puzzles as you can.

Instructor note: the next pages consist of one puzzle to solve per page, which is replaced by a new puzzle when solved, until a four minute timer expires.

Page 6. Please complete the following short survey.

What is your age?

What is your gender?

_ Male

_ Female

_ Other

Which would you say more closely describes your racial or ethnic background?

_ White

_ Black

_ Asian or Pacific Islander

_ Hispanic

_ Native-American

_ Other

What is the highest degree or level of schooling you have completed? If currently enrolled, highest degree received.

In which US state do you reside?

In which city do you reside?

Are you currently employed outside of Mechanical Turk? If so, what is your job? For how many months have you been working on Mechanical Turk?

Please list your three favorite hobbies:

Page 7. You have earned *'puzzle payment'* and your total earnings including your bonus equal *'earnings'*.

Please press next to finish the study.

B.2.2 Employer Treatment B

Page 1. There are two roles in this study: workers and employers. You will play the role of the employer. Participants assigned to be workers were previously tasked with solving as many computer puzzles as they could in a 4 minute period and their performance was recorded. They were paid 250 credits for each puzzle solved.

As an employer, your task is to identify and select the workers who have solved the most puzzles from the pool of available workers over 15 periods of play. In each period, you will select one worker and learn how many puzzles they solved. Each worker is employed for only one period. Below, you are shown an example of a puzzle similar to those solved by workers. The square with characters on the right differs from the left in two letters. Workers had to identify those letters to solve the puzzle.

m	l	h	a	t	o
s	p	t	q	g	b
u	p	b	c	c	w
v	i	l	e	e	j
x	v	t	l	l	m
u	y	e	x	k	a
y	z	l	u	b	i

m	l	h	e	t	o
s	p	t	q	g	b
u	p	b	l	c	w
v	i	l	e	e	j
x	v	t	l	l	m
u	y	e	x	k	a
y	z	l	u	b	i

Page 2. You are given the following information about workers. There are 200 workers in total from two distinct groups, group Orange and group Gray. Group Orange is composed of 75% of all workers (150 workers) while group Gray is composed of 25% of all workers (50 workers). When you select a worker from the Orange group, you will automatically be given the group average of 9 puzzles solved. No initial information is given on how many puzzles Gray workers solved on average, but you can learn about the group by hiring workers from it.

When you select a worker, the number of puzzles they solved will be revealed to you and you will receive 220 credits for each puzzle that they were able to solve. For example, whenever you hire from group Orange, the number of puzzles will be 9 and you will earn 9 x 220 = 1980 credits. You will be paid based on credits earned in 5 of the 15 periods of play, which will be randomly selected by the computer at the end of the study. In every period, maximizing expected earnings therefore requires selecting workers from the group which can solve more puzzles.

Page 3. Before you begin selecting workers and after every period in which you hire a Gray worker, you will be asked what you believe the average number of puzzles solved by Gray workers was given the information available to you. If you hire from group Orange in a given period, your prediction about group Gray will carry over from the last time you had to predict the number of puzzles they solved.

You will receive a small number of extra credits for having more accurate predictions in two periods. These two periods will be selected randomly by the computer at the end of the study.

A procedure was designed to ensure that it is always in your interest to give your best prediction of the average number of puzzles solved by group Gray. The details are as follows. When you make your prediction, the difference between the true productivity of the group and your prediction is calculated and then squared. This number constitutes your prediction error and is compared to a random number (not shown) drawn between 0 and 81 where each number has an equal probability of being drawn.

If that random number is greater or equal to your prediction error and the period was selected for payment, you will receive 110 extra credits for the period. Otherwise, you will receive 0 extra credits.

You will be shown a few examples on the next page to familiarize yourself with the procedure.

Page 4. Here are a few examples of how your prediction earnings are determined.

Suppose that workers from group Gray were able to solve 7 puzzles on average and that the random number drawn for the period was 4.

If your prediction is 8, the difference between your prediction and the true average is 8 - 7 = 1. Your prediction error is 1x1 = 1. Since 1 is no greater than 4, you receive 110 credits.

Suppose instead that workers from group Gray were able to solve 10 puzzles on average and that the random number drawn for the period was 8.

If your prediction is 8, the difference between your prediction and the true average is 8 - 10 = -2. Your prediction error is -2x-2 = 4. Since 4 is no greater than 8, you receive 110 credits.

If your prediction is 13, the difference between your prediction and the true average is 13 - 10 = 3. Your prediction error is 3x3 = 9. Since 9 is greater than 8, you receive 0 credits.

As such, making a prediction which is closer to the truth increases your chance of receiving extra credits.

Page 5. Here are two last important points you should keep in mind. Credits from hiring a worker who solved one more puzzle in a single period (220 per puzzle) are equal to the total extra credits which you can earn from making better predictions about group Gray (220 in total) over all periods. As such, maximizing expected credits requires hiring workers who solved more puzzles.

Similarly, since the majority of your earnings will be based on your hiring performance rather than participation, maximizing payment also requires maximizing the number of puzzles solved by your employees.

Here is a summary of other key points.

Over 15 periods, you will select one worker per period from either group Gray or group Orange and observe the number of puzzles they solved.

Group Orange has 150 workers, group Gray has 50 workers. On average, group Orange workers solved 9 puzzles, which is the value that will be given to you if you select from that group.

In five of these periods, you will be paid 220 credits for each puzzle that the workers you select were able to solve. Before the first period and after every period in which you hire a Gray worker, you will be asked about your best prediction of the average number of puzzles solved by Gray workers.

In two periods, you will receive 110 extra credits if your current prediction is close enough to the truth. On the next page, the instructions will be repeated in their entirety for your reference.

Page 6. The entire instructions are repeated below for reference. Please review them as needed. Throughout the study, you cannot return to previous pages once you advance to a new page. To ensure that you understand the instructions before proceeding, some comprehension questions will be asked at the bottom of this page.

Instructor note: the entire instructions were repeated here.

Please answer the following comprehension questions. If you make a mistake, please refer to the instructions above. You will not be able to advance to the study until you answer all questions correctly.

1. What is your role in this study? Enter 1 for "Employer" and 2 for "Worker".

2. For how many periods will you hire workers?

3. Which group has the most workers? Enter 1 for "Gray", 2 for "Orange", 3 for "Same" and 4 for "Unknown".'

4. How many workers can you hire per period?

5. On average, how many puzzles did group Orange solve?

6. You will receive credits for every puzzle solved by your workers. How many credits will you receive per puzzle?

7. On average, did group Gray solve fewer or more puzzles than group Orange? Enter 1 for "Fewer", 2 for "Same", 3 for "More" and 4 for "Unknown".

8. How many credits in total can you earn from accurately predicting the average number of puzzles solved by group Gray across all periods?

9. Suppose group Gray solved 10 puzzles on average and your prediction for the average number of puzzles solved by the group is 12. Your prediction error is then 4. What is the smallest random number that can be drawn and still earn you credits?

10. Which is worth the most credits overall, predicting the average number of puzzles solved by group Gray or hiring better workers from any group? Enter 1 for "Predicting", 2 for "Hiring", 3 for "Same" and 4 for "Unknown".

Page 7. Given the information available to you so far, report your best prediction of the average number of puzzles solved by group Gray:

Page 8. Period 1 of 15.

Please choose the group from which you want to select a worker:

Page 9. Period 1 of 15.

The worker you have selected solved *'no. puzzle'* puzzles and you have earned *'no. credit'* credits.

Instructor note: the following is only displayed if a worker from group Gray was hired.

Given the information available to you so far, report your best prediction of the average number of puzzles solved by group Gray:

Note that your last prediction was *'prior'*.

Instructor note: pages 8 and 9 are repeated for periods 2-15.

Page 10. Please imagine the following hypothetical scenario. There are two urns, A and B, with balls of color red and black. Urn A has 50 Red balls and 50 Black balls. Urn B has an unknown number of Red and Black balls (with a total of 100 balls).

You will first select a color, Red or Black, and this will be your Success Color. One ball would then be drawn from one of the urns.

Urn A pays 200 tokens if the ball is drawn from it and matches your Success Color, and 0 tokens if it does not match. Since each color has a 1/2 chance of being drawn, this means that drawing from Urn A pays 200 tokens with a chance of 1/2, and pays 0 with a chance of 1/2.

Urn B pays a positive amount if the ball drawn from it matches your Success Color, and 0 tokens if it does not match. Since the chance of each color being drawn is unknown, the chance of Urn B paying a positive number of tokens is unknown as well.

The table below shows 20 cases which increase in the amount paid when a ball matching your Success Color is drawn from Urn B. One of these cases will be selected for payment at random. Your task is to choose at which case (number between 1 and 20) you want to "Switch" from drawing from Urn A to Urn B.

Making a choice to switch means that for every case before your choice, a ball would be drawn from Urn A. For each case after your choice, including the case for which you switch, a ball would be drawn from Urn B.

For example, if the case randomly selected is 9, case No. 9 would determine payment.

- If your "Switch" number is higher than 9, a ball would be drawn from Urn A, and if the color of the ball matches the chosen Success Color, then you would earn 200 tokens. If it does not match, you would earn 0 tokens.

- If your "Switch" number is 9 or lower, a ball would be drawn from Urn B, and if the color of the ball matches the chosen Success Color, then you would earn 228 tokens. If it does not match, you would earn 0 tokens.

Please make your choices by selecting from the drop-down lists below the table.

	Urn A 50 Red Balls, 50 Black Balls	Urn B ? Red Balls, ? Black Balls
1	200 tokens if Chosen Color, 0 tokens if not	164 tokens if Chosen Color, 0 tokens if not
2	200 tokens if Chosen Color, 0 tokens if not	172 tokens if Chosen Color, 0 tokens if not
3	200 tokens if Chosen Color, 0 tokens if not	180 tokens if Chosen Color, 0 tokens if not
4	200 tokens if Chosen Color, 0 tokens if not	188 tokens if Chosen Color, 0 tokens if not
5	200 tokens if Chosen Color, 0 tokens if not	196 tokens if Chosen Color, 0 tokens if not
6	200 tokens if Chosen Color, 0 tokens if not	204 tokens if Chosen Color, 0 tokens if not
7	200 tokens if Chosen Color, 0 tokens if not	212 tokens if Chosen Color, 0 tokens if not
8	200 tokens if Chosen Color, 0 tokens if not	220 tokens if Chosen Color, 0 tokens if not
9	200 tokens if Chosen Color, 0 tokens if not	228 tokens if Chosen Color, 0 tokens if not
10	200 tokens if Chosen Color, 0 tokens if not	236 tokens if Chosen Color, 0 tokens if not
11	200 tokens if Chosen Color, 0 tokens if not	244 tokens if Chosen Color, 0 tokens if not
12	200 tokens if Chosen Color, 0 tokens if not	252 tokens if Chosen Color, 0 tokens if not
13	200 tokens if Chosen Color, 0 tokens if not	260 tokens if Chosen Color, 0 tokens if not
14	200 tokens if Chosen Color, 0 tokens if not	268 tokens if Chosen Color, 0 tokens if not
15	200 tokens if Chosen Color, 0 tokens if not	276 tokens if Chosen Color, 0 tokens if not
16	200 tokens if Chosen Color, 0 tokens if not	284 tokens if Chosen Color, 0 tokens if not
17	200 tokens if Chosen Color, 0 tokens if not	292 tokens if Chosen Color, 0 tokens if not
18	200 tokens if Chosen Color, 0 tokens if not	300 tokens if Chosen Color, 0 tokens if not
19	200 tokens if Chosen Color, 0 tokens if not	308 tokens if Chosen Color, 0 tokens if not
20	200 tokens if Chosen Color, 0 tokens if not	316 tokens if Chosen Color, 0 tokens if not

Page 11. You have completed the main tasks of the study, thank you. A short survey will now follow. Completion of this survey is also required for payment.

Page 12. Please complete the following short survey (1 of 3). The following questions relate to your beliefs about the groups and your selection decisions.

Do you believe group Gray solved fewer, the same, or more puzzles than group Orange on average?

How important do you think intelligence is to explain the difference?

Very Important, Important, Somewhat Important, Unimportant, Very Unimportant

How important do you think effort is to explain this difference?

Very Important, Important, Somewhat Important, Unimportant, Very Unimportant

How important do you think experience or practice is to explain this difference?

Very Important, Important, Somewhat Important, Unimportant, Very Unimportant

Selecting from group Gray was riskier than selecting from group Orange.

Very Important, Important, Somewhat Important, Unimportant, Very Unimportant

Page 13. Please complete the following short survey (2 of 3). The following questions ask some information about yourself.

What is your age?

What is your gender?

_ Male

_ Female

_ Other

Which would you say more closely describes your racial or ethnic background?

_ White

_ Black

_ Asian or Pacific Islander

_ Hispanic

_ Native-American

_ Other

What is the highest degree or level of schooling you have completed? If currently enrolled, highest degree received.

In which US state do you reside?

In which city do you reside?

Are you currently employed outside of Mechanical Turk? If so, what is your job?

Page 14. Please complete the following short survey (3 of 3). The following questions ask about some of your views on race. Report the extent to which you agree with the following statements: (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

You oppose policies which give a preference in hiring and promotion to African-Americans.

In general, African-Americans are as hard-working as whites.

In general, African-Americans are as competent at their job as whites.

In general, African-Americans are as intelligent as whites.

In general, African-Americans have as much schooling as whites.

There should be laws against marriages between African-Americans and whites.

If you are African-American, please consider the following questions as relating to whites rather than African-Americans.

Are any members of your family or close friends African-American?

In general, you feel close to African-Americans. (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

You would object if a family member brought an African-American friend home for dinner. (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

If an African-American with the same income and education as you have moved in to your block, this would make a difference to you. (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

You object sending your children to a school with more than a few African-American students. (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

If your political party nominated an African-American for president, you would vote for them if they were qualified for the job. (Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree)

Page 15. Thank you for participating.

You have earned *'task payment'*. Your total earnings including both your participation prize and bonus equal *'earnings'*.

Please press next to finish the study. You must press next to guarantee your payment.

B.2.3 Employer Treatment C (Only pages with differences from Treatment B)

Page 2. You are given the following information about workers. There are 200 workers in total from two distinct groups, group Orange and group Gray. Group Orange is composed of 75% of all workers (150 workers) while group Gray is composed of 25% of all workers (50 workers).

In this study, you will focus on hiring from the group of 50 Gray workers. No initial information is given on how many puzzles Gray workers solved on average, but you can learn about the group by hiring workers from it. Those from group Orange were able to solve 9 puzzles on average.

When you select a worker, the number of puzzles they solved will be revealed to you and you will receive 180 credits for each puzzle that they were able to solve.

You will be paid based on credits earned in 5 of the 15 periods of play, which will be randomly selected by the computer at the end of the study.

Page 3. Before you begin selecting workers and after every period, you will be asked what you believe the average number of puzzles solved by Gray workers was given the information available to you.

You will receive a small number of extra credits for having more accurate predictions in two periods. These two periods will be selected randomly by the computer at the end of the study.

A procedure was designed to ensure that it is always in your interest to give your best prediction of the average number of puzzles solved by workers. The details are as follows. When you make your prediction, the difference between the true productivity of the group and your prediction is calculated and then squared. This number constitutes your prediction error and is compared to a random number (not shown) drawn between 0 and 81 where each number has an equal probability of being drawn.

If that random number is greater or equal to your prediction error and the period was selected for payment, you will receive 110 extra credits for the period. Otherwise, you will receive 0 extra credits.

You will be shown a few examples on the next page to familiarize yourself with the procedure.

Page 4. Here are a few examples of how your prediction earnings are determined.

Suppose that workers were able to solve 7 puzzles on average and that the random number drawn for the period was 4. If your prediction is 8, the difference between your prediction and the true average is 8 - 7 = 1. Your prediction error is 1x1 = 1. Since 1 is no greater than 4, you receive 110 credits.

Suppose instead that workers were able to solve 10 puzzles on average and that the random number drawn for the period was 8.

If your prediction is 8, the difference between your prediction and the true average is 8 - 10 = -2. Your prediction error is -2x-2 = 4. Since 4 is no greater than 8, you receive 110 credits.

If your prediction is 13, the difference between your prediction and the true average is 13 - 10 = 3. Your prediction error is 3x3 = 9. Since 9 is greater than 8, you receive 0 credits.

As such, making a prediction which is closer to the truth increases your chance of receiving extra credits.

Page 5. Here are two last important points you should keep in mind: Credits from hiring a worker who solved one more puzzle in a single period (180 per puzzle) are equal to the total extra credits which you can earn from making better predictions over all periods (also 180).

Since the majority of your earnings will be based on your hiring performance rather than participation, payment mostly depends on the number of puzzles solved by your workers.

Here is a summary of other key points.

Over 15 periods, you will select one of the 50 Gray workers per period and observe the number of puzzles they solved. In five of these periods, you will be paid 180 credits for each puzzle that the workers you select were able to solve. Before the first period and after every period, you will be asked about your best prediction of the average number of puzzles solved by Gray workers.

In two periods, you will receive 90 extra credits if your current prediction is close enough to the truth.

On the next page, the instructions will be repeated in their entirety for your reference.

Page 6. The entire instructions are repeated below for reference. Please review them as needed. Throughout the study, you cannot return to previous pages once you advance to a new page. To ensure that you understand the instructions before proceeding, some comprehension questions will be asked at the bottom of this page.

Instructor note: the entire instructions were repeated here.

Please answer the following comprehension questions. If you make a mistake, please refer to the instructions above. You will not be able to advance to the study until you answer all questions correctly.

1. What is your role in this study? Enter 1 for "Employer" and 2 for "Worker".

2. For how many periods will you hire workers?

3. From which group of workers are you hiring? Enter 1 for "Orange" and 2 for "Gray".

4. How many workers can you hire per period?

5. You will receive credits for every puzzle solved by your workers. How many credits will you receive per puzzle?

6. How many credits in total can you earn from accurately predicting the average number of puzzles solved by workers across all periods?

7. Suppose workers solved 10 puzzles on average and your prediction for the average number of puzzles is 12. Your prediction error is then 4. What is the smallest random number that can be drawn and still earn you credits?

8. Which is worth the most credits overall, predicting the average number of puzzles solved by workers or hiring better workers? Enter 1 for "Predicting", 2 for "Hiring", 3 for "Same" and 4 for "Unknown".

9. How many workers are there in group Gray?

10. How many puzzles did workers from group Orange solve on average?

Page 7. Given the information available to you so far, report your best prediction of the average number of puzzles solved by Gray workers:

Page 8. Period 1 of 15.

Please hire a worker by selecting below:

Page 9. Period 1 of 15.

The worker you have selected solved *'no puzzle'* puzzles and you have earned *'no credit'* credits.

Given the information available to you so far, report your best prediction of the average number of puzzles solved by Gray workers:

Note that your last prediction was *'prior'*.

Instructor note: pages 8 and 9 are repeated for periods 2-15.

Page 10. Please complete the following short survey.

What is your age?

What is your gender?

_ Male

_ Female

_ Other

Which would you say more closely describes your racial or ethnic background?

_ White

_ Black

_ Asian or Pacific Islander

_ Hispanic

_ Native-American

_ Other

What is the highest degree or level of schooling you have completed? If currently enrolled, highest degree received.

In which US state do you reside?

In which city do you reside?

Are you currently employed outside of Mechanical Turk? If so, what is your job?

Page 11. Thank you for participating.

You have earned *'task payment'*. Your total earnings including both your participation prize and bonus equal *'earnings'*.

Please press next to finish the study. You must press next to guarantee your payment.

B.2.4 Employer Treatment B1 (Only pages with differences from Treatment B)

Page 2. You are given the following information about workers. There are 100 workers in total from two distinct groups, group Gray and group Orange. There are 50 workers in each group.

When you select a worker from the Orange group, you will automatically be given the group average of 9 puzzles solved.

No initial information is given on how many puzzles Gray workers solved on average, but you can learn about the group by hiring workers from it.

When you select a worker, the number of puzzles they solved will be revealed to you and you will receive 220 credits for each puzzle that they were able to solve. For example, whenever you hire from group Orange, the number of puzzles will be 9 and you will earn 9 x 220 = 1980 credits.

You will be paid based on credits earned in 5 of the 15 periods of play, which will be randomly selected by the computer at the end of the study. In every period, maximizing expected earnings therefore requires selecting workers from the group which can solve more puzzles.

Page 5. Here are two last important points you should keep in mind. Credits from hiring a worker who solved one more puzzle in a single period (220 per puzzle) are equal to the total extra credits which you can earn from making better predictions about group Gray (220 in total) over all periods. As such, maximizing expected credits requires hiring workers who solved more puzzles.

Similarly, since the majority of your earnings will be based on your hiring performance rather than participation, maximizing payment also requires maximizing the number of puzzles solved by your employees.

Here is a summary of other key points.

Over 15 periods, you will select one worker per period from either group Gray or group Orange and observe the number of puzzles they solved.

Both groups have 50 workers. On average, group Orange workers solved 9 puzzles, which is the value that will be given to you if you select from that group.

In five of these periods, you will be paid 220 credits for each puzzle that the workers you select were able to solve. Before the first period and after every period in which you hire a Gray worker, you will be asked about your best prediction of the average number of puzzles solved by Gray workers.

In two periods, you will receive 110 extra credits if your current prediction is close enough to the truth.

On the next page, the instructions will be repeated in their entirety for your reference.

B.2.5 Employer Treatment B2 (Only pages with differences from Treatment B)

Page 3. The entire instructions are repeated below for reference. Please review them as needed. Throughout the study, you cannot return to previous pages once you advance to a new page. To ensure that you understand the instructions before proceeding, some comprehension questions will be asked at the bottom of this page.

Instructor note: the entire instructions were repeated here.

Please answer the following comprehension questions. If you make a mistake, please refer to the instructions above. You will not be able to advance to the study until you answer all questions correctly.

1. What is your role in this study? Enter 1 for "Employer" and 2 for "Worker".

2. For how many periods will you hire workers?

3. Which group has the most workers? Enter 1 for "Gray", 2 for "Orange", 3 for "Same" and 4 for "Unknown".

4. How many workers can you hire per period?

5. On average, how many puzzles did group Orange solve?

6. You will receive credits for every puzzle solved by your workers. How many credits will you receive per puzzle?

7. On average, did group Gray solve fewer or more puzzles than group Orange? Enter 1 for "Fewer", 2 for "Same", 3 for "More" and 4 for "Unknown".

Page 4. Period 1 of 15.

Please choose the group from which you want to select a worker:

Page 5. Period 1 of 15.

The worker you have selected solved *'no puzzle'* puzzles and you have earned *'no credit'* credits.

Instructor note: pages 4 and 5 are repeated for periods 2-15.

Page 6. You will now be asked what you believe the average number of puzzles solved by Gray workers was given the information available to you. You will receive extra credits for a more accurate prediction.

A procedure was designed to ensure that it is always in your interest to give your best prediction of the average number of puzzles solved by group Gray. The details are as follows. When you make your prediction, the difference between the true productivity of the group and your prediction is calculated and then squared. This number constitutes your prediction error and is compared to a random number (not shown) drawn between 0 and 81 where each number has an equal probability of being drawn.

If that random number is greater or equal to your prediction error, you will receive 220 extra credits. Otherwise, you will receive 0 extra credits.

Here are a few examples of how your prediction earnings are determined.

Suppose that workers from group Gray were able to solve 7 puzzles on average and that the random number drawn was 4. If your prediction is 8, the difference between your prediction and the true average is 8 - 7 = 1. Your prediction error is 1x1 = 1. Since 1 is no greater than 4, you receive 220 credits.

Suppose instead that workers from group Gray were able to solve 10 puzzles on average and that the random number drawn was 8.

If your prediction is 8, the difference between your prediction and the true average is 8 - 10 = -2. Your prediction error is -2x-2 = 4. Since 4 is no greater than 8, you receive 220 credits.

If your prediction is 13, the difference between your prediction and the true average is 13 - 10 = 3. Your prediction error is 3x3 = 9. Since 9 is greater than 8, you receive 0 credits.

As such, making a prediction which is closer to the truth increases your chance of receiving extra credits. Please answer the following comprehension question. If you make a mistake, please refer to the instructions above. You will not be able to advance until you answer the question correctly.

Suppose group Gray solved 10 puzzles on average and your prediction for the average number of puzzles solved by the group is 12. Your prediction error is then 4. What is the smallest random number that can be drawn and still earn you credits?

Page 7. Given the information available to you so far, report your best prediction of the average number of puzzles solved by group Gray:

Instructor note: pages 8-11 correspond to Treatment B pages 12-15.

APPENDIX C

Appendix for Chapter IV

C.1 Additional Information on Hiring Experiences

Table C.1: Hiring Experiences and Worker Performance

	Bad Experience	Good Experience
Promotion	0	0.002
		(0.045)
Sale Performance	-0.138	-0.018
	(0.572)	(0.033)
Termination Reason		
Job Abandonment	37.78	17.28
Probationary Period Termination	4.69	0.13
Unsatisfactory Performance	4.65	0.33
Career Advancement	2.66	9.60
Personal Reasons	10.30	18.66
Return to school	5.08	9.70

Sale performance is a normalized measure of sale performance relative to targets established by the firm.

C.2 More on Heterogeneity in Group Hiring Across Managers

Figure C.1: Manager Predicted Shares of White Hiring Based on Hiring Context, Market Factors, and Manager Fixed Effects

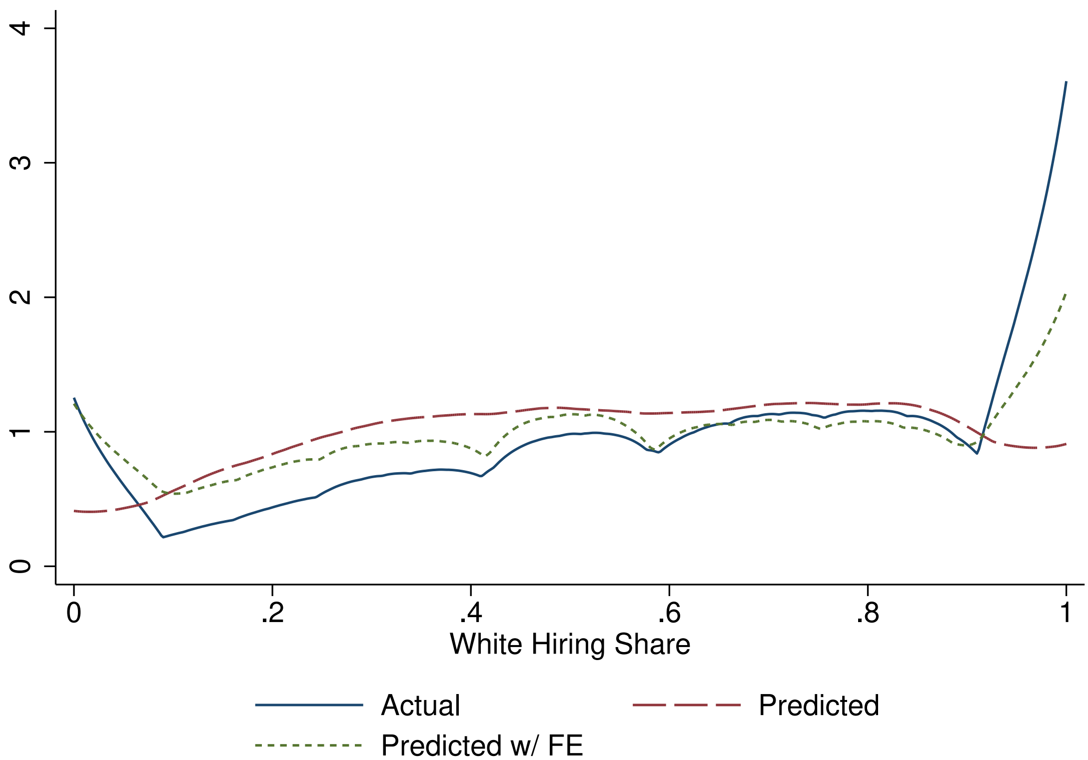

See Figure 4.2 for details.

Figure C.2: Distribution of Manager Fixed Effects for White Hiring

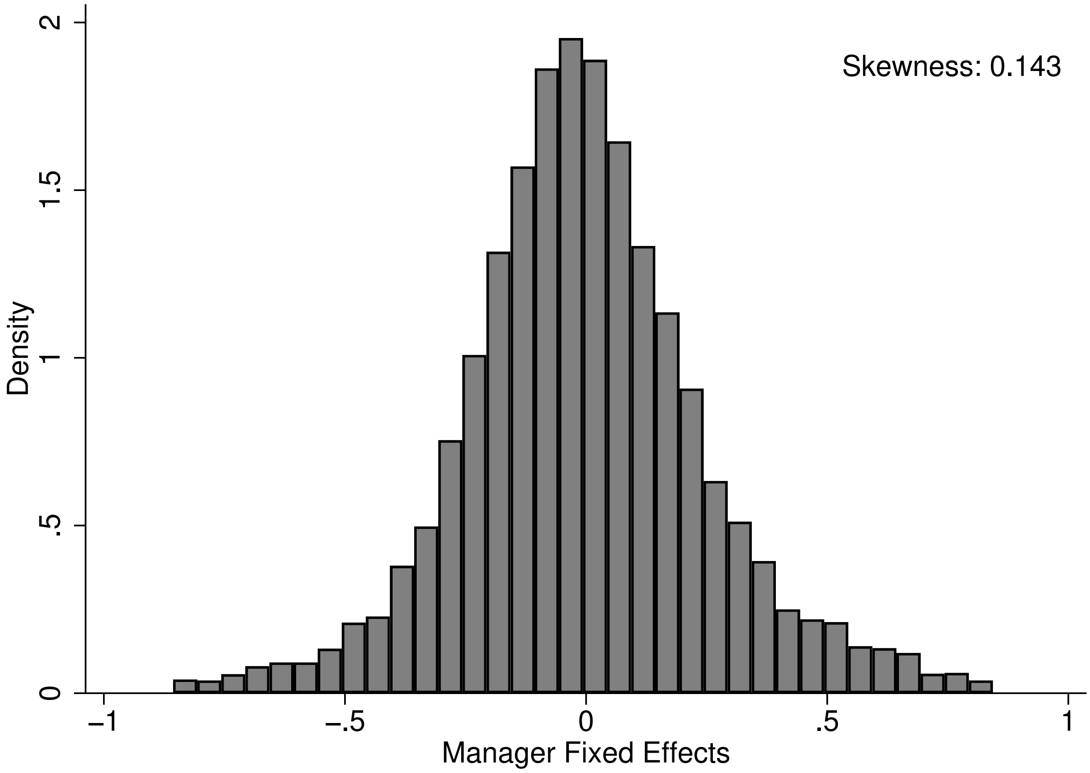

See Figures 4.2 and 4.3 for details.

Figure C.3: Manager Predicted Shares of Black Hiring Based on Hiring Context, Market Factors, and Manager Fixed Effects, Restricted to Managers with at least 5 Hires

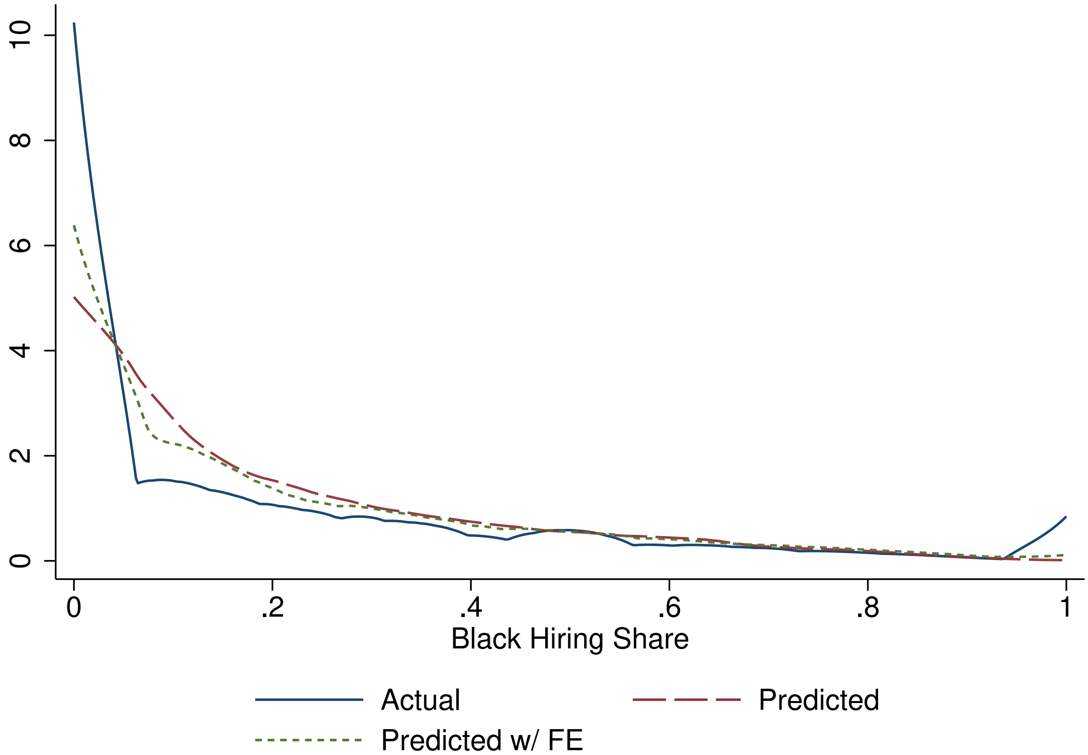

See Figure 4.2 for details.

C.3 Additional Results on the Impact of Hiring Experiences

Table C.2: Robustness Checks and Heterogeneity Analyses

	Fired	Quit	Relative to White	Relative to CBSA	High Black Population	Low Black Population	"Exogenous" Separation	"Endowed" Workers	Store Change	Avg. Exp. versus Others'
Cumulative Average										
Fraction Quit or Fired										
Black	-0.087	-0.057	-0.068	-0.079	-0.091	-0.061	-0.063	-0.034	-0.067	-0.069
	(0.027)	(0.018)	(0.014)	(0.016)	(0.022)	(0.023)	(0.032)	(0.027)	(0.045)	(0.023)
White	0.112	0.025								
	(0.051)	(0.024)								
Other Managers at the Store										
Black										-0.016
										(0.025)
White										-0.004
										(0.037)
Outcome Mean	0.367	0.37	0.348	0.369	0.470	0.181	0.37	0.328	0.356	0.352
Observations	33,971	33,971	31,911	33,675	21,786	12,140	33,971	10,911	977	30,985
Fraction Long Tenure										
Black			0.033	0.031	0.074	0.025		-0.006	0.010	0.057
			(0.021)	(0.024)	(0.029)	(0.034)		(0.026)	(0.074)	(0.034)
Other Managers at the Store										
Black										0.007
										(0.037)
White										0.025
										(0.028)
Outcome Mean			0.37	0.37	0.47	0.18		0.33	0.35	0.35
Observations			28,456	28,456	18,354	10,074		9,360	790	25,916

Clustered standard errors at the manager level are presented in parentheses. High (Low) black population refer to the store being located in a CBSA with above (below) median black population. Exogenous Separation restricts fires and quits to dissatisfaction with pay, compensation or benefits, worker integrity, illegal or unethical behavior, or violation of rules and policies. Endowed Workers corresponds to workers already in the department at the manager's arrival. See Table 4.2 for additional details.

C.4 Hispanic Workers

Table C.3: Cumulative Impact of Previous Experiences with Hispanic and White Workers on Current Hiring of Hispanic Workers

Expected Tenure	
Hispanic	0.023
	(0.015)
White	0.003
	(0.013)
Outcome Mean	0.320
Observations	29,378
Fraction Quit or Fired	
Hispanic	-0.025
	(0.019)
White	0.029
	(0.023)
Outcome Mean	0.293
Observations	27,349
Fraction Long Tenure	
Hispanic	0.003
	(0.024)
White	-0.016
	(0.027)
Outcome Mean	0.290
Observations	22,482

See Table 4.2 for details.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Abowd, J.M., Kramarz, F. and Margolis, D.N., 1999. High wage workers and high wage firms. Econometrica, 67(2), pp.251-333.

[2] Agnew, J.R., Bateman, H., Eckert, C., Iskhakov, F., Louviere, J. and Thorp, S., 2018. First impressions matter: An experimental investigation of online financial advice. Management Science, 64(1), pp.288-307.

[3] Aguirregabiria, V. and Jeon, J., 2019. Firms' beliefs and learning: Models, identification, and empirical evidence. Review of Industrial Organization, 56, pp.203-235.

[4] Aigner, D.J. and Cain, G.G., 1977. Statistical theories of discrimination in labor markets. ILR Review, 30(2), pp.175-187.

[5] Altonji, J.G. and Pierret, C.R., 2001. Employer learning and statistical discrimination. The Quarterly Journal of Economics, 116(1), pp.313-350.

[6] Allport, G.W., 1954. The nature of prejudice. Oxford, England: Addison-Wesley

[7] Anderson, C.M., 2001. Behavioral models of strategies in multi-armed bandit problems. Doctoral dissertation, California Institute of Technology.

[8] Anderson, C.M., 2012. Ambiguity aversion in multi-armed bandit problems. Theory and Decision, 72(1), pp.15-33.

[9] Arcidiacono, P., Bayer, P. and Hizmo, A., 2010. Beyond signaling and human capital: Education and the revelation of ability. American Economic Journal: Applied Economics, 2(4), pp.76-104.

[10] Arnold, D., Dobbie, W. and Yang, C.S., 2018. Racial bias in bail decisions. The Quarterly Journal of Economics, 133(4), pp.1885-1932.

[11] Arrow, K., 1973. The theory of discrimination. Discrimination in Labor Markets, 3(10), pp.3-33.

[12] Ashenfelter, O. and Hannan, T., 1986. Sex discrimination and product market

competition: The case of the banking industry. The Quarterly Journal of Economics, 101(1), pp.149-173.

[13] Audretsch, D.B., 1991. New-firm survival and the technological regime. The Review of Economics and Statistics, 73(3), pp.441-450.

[14] Autor, D.H. and Scarborough, D., 2008. Does job testing harm minority workers? Evidence from retail establishments. The Quarterly Journal of Economics, 123(1), pp.219-277.

[15] Baliga, S., Hanany, E. and Klibanoff, P., 2013. Polarization and ambiguity. American Economic Review, 103(7), pp.3071-83.

[16] Banerjee, S., 2011. Learning from prices and the dispersion in beliefs. The Review of Financial Studies, 24(9), pp.3025-3068.

[17] Banerjee, S., Kaniel, R. and Kremer, I., 2009. Price drift as an outcome of differences in higher-order beliefs. The Review of Financial Studies, 22(9), pp.3707-3734.

[18] Banerjee, S. and Kremer, I., 2010. Disagreement and learning: Dynamic patterns of trade. The Journal of Finance, 65(4), pp.1269-1302.

[19] Banks, J., Olson, M. and Porter, D., 1997. An experimental analysis of the bandit problem. Economic Theory, 10(1), pp.55-77.

[20] Bardhi, A., Guo, Y. and Strulovici, B., 2020. Early-Career Discrimination: Spiraling or Self-Correcting?. Working Paper. Northwestern University.

[21] Bartoš, V., Bauer, M., Chytilová, J. and Matějka, F., 2016. Attention discrimination: Theory and field experiments with monitoring information acquisition. American Economic Review, 106(6), pp.1437-75.

[22] Becker, G.S., 1957. The economics of discrimination. University of Chicago press.

[23] Benson, A., Board, S. and Meyer-ter-Vehn, M., 2019. Discrimination in hiring: Evidence from retail sales. Unpublished. University of Minnesota.

[24] Benson, A., Li, D. and Shue, K., 2019. Promotions and the Peter principle. The Quarterly Journal of Economics, 134(4), pp.2085-2134.

[25] Bentley, F.R., Daskalova, N. and White, B., 2017, May. Comparing the reliability of Amazon Mechanical Turk and Survey Monkey to traditional market research surveys. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 1092-1099.

[26] Bergemann, D. and Välimäki, J., 2008. Bandit problems. The New Palgrave Dictionary of Economics: Volume 1–8, pp.336-340.

[27] Bergman, P., Li, D. and Raymond, L., 2020. Hiring as Exploration. SSRN working paper 3630630.

[28] Berinsky, A.J., Huber, G.A. and Lenz, G.S., 2012. Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk. Political Analysis, 20(3), pp.351-368.

[29] Berson, C., Laouénan, M. and Valat, E., 2019. Outsourcing recruitment as a solution to prevent discrimination: A Correspondence Study. IZA Discussion Paper No. 12132.

[30] Bertrand, M. and Mullainathan, S., 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. American Economic Review, 94(4), pp.991-1013.

[31] Bertrand, M., Chugh, D. and Mullainathan, S., 2005. Implicit discrimination. American Economic Review, 95(2), pp.94-98.

[32] Bertrand, M. and Duflo, E., 2017. Field experiments on discrimination. In Handbook of Economic Field Experiments (Vol. 1, pp. 309-393). North-Holland.

[33] Biernat, M. and Kobrynowicz, D., 1997. Gender-and race-based standards of competence: lower minimum standards but higher ability standards for devalued groups. Journal of Personality and Social Psychology, 72(3), p.544.

[34] Black, D.A., 1995. Discrimination in an equilibrium search model. Journal of Labor Economics, 13(2), pp.309-334.

[35] Black, S.E. and Strahan, P.E., 2001. The division of spoils: rent-sharing and discrimination in a regulated industry. American Economic Review, 91(4), pp.814-831.

[36] Blau, F.D. and Kahn, L.M., 2017. The gender wage gap: Extent, trends, and explanations. Journal of Economic Literature, 55(3), pp.789-865.

[37] Bohren, J.A., Haggag, K., Imas, A. and Pope, D.G., 2019a. Inaccurate statistical discrimination (No. w25935). National Bureau of Economic Research.

[38] Bohren, J.A., Imas, A. and Rosenberg, M., 2019b. The dynamics of discrimination: Theory and evidence. American Economic Review, 109(10), pp.3395-3436.

[39] Bordalo, P., Coffman, K., Gennaioli, N. and Shleifer, A., 2016. Stereotypes. The Quarterly Journal of Economics, 131(4), pp.1753-1794.

[40] Bordalo, P., Coffman, K., Gennaioli, N. and Shleifer, A., 2019. Beliefs about gender. American Economic Review, 109(3), pp.739-73.

[41] Buhrmester, M., Kwang, T. and Gosling, S.D., 2011. Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?. Perspectives on Psychological Science, 6(1), pp.3-5.

[42] Card, D., Cardoso, A.R. and Kline, P., 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. The Quarterly Journal of Economics, 131(2), pp.633-686.

[43] Charles, K.K. and Guryan, J., 2011. Studying discrimination: Fundamental challenges and recent progress. Annu. Rev. Econ., 3(1), pp.479-511.

[44] Charness, G. and Kuhn, P., 2011. Lab labor: What can labor economists learn from the lab?. In Handbook of Labor Economics (Vol. 4, pp. 229-330). Elsevier.

[45] Charness, G., Rigotti, L. and Rustichini, A., 2007. Individual behavior and group membership. American Economic Review, 97(4), pp.1340-1352.

[46] Che, Y.K., Kim, K. and Zhong, W., 2020. Statistical discrimination in ratings-guided markets. arXiv preprint arXiv:2004.11531.

[47] Chen, Y. and Li, S.X., 2009. Group identity and social preferences. American Economic Review, 99(1), pp.431-57.

[48] Chen, D.L., Schonger, M. and Wickens, C., 2016. oTree — An open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance, 9, pp.88-97.

[49] Coate, S. and Loury, G.C., 1993. Will affirmative-action policies eliminate negative stereotypes?. American Economic Review, 83(5), pp.1220-1240.

[50] Cornell, B. and Welch, I., 1996. Culture, information, and screening discrimination. Journal of Political Economy, 104(3), pp.542-571.

[51] Coutts, A., 2019. Good news and bad news are still news: Experimental evidence on belief updating. Experimental Economics, 22(2), pp.369-395.

[52] DeGroot, M.H., 1974. Reaching a consensus. Journal of the American Statistical Association, 69(345), pp.118-121.

[53] De Long, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of political Economy, 98(4), pp.703-738.

[54] Dianat, A., Echenique, F. and Yariv, L., 2018. Statistical discrimination and affirmative action in the lab. CEPR Discussion Paper No. DP12915.

[55] Enke, B. and Zimmermann, F., 2019. Correlation neglect in belief formation. The Review of Economic Studies, 86(1), pp.313-332.

[56] Eyster, E. and Piccione, M., 2013. An approach to asset pricing under incomplete and diverse perceptions. Econometrica, 81(4), pp.1483-1506.

[57] Eyster, E., Rabin, M. and Vayanos, D., 2019. Financial markets where traders neglect the informational content of prices. The Journal of Finance, 74(1), pp.371-399.

[58] Fang, H. and Moro, A., 2011. Theories of statistical discrimination and affirmative action: A survey. In Handbook of Social Economics (Vol. 1, pp. 133-200). North-Holland.

[59] Farber, H.S. and Gibbons, R., 1996. Learning and wage dynamics. The Quarterly Journal of Economics, 111(4), pp.1007-1047.

[60] Farrell, D., Wheat, C. and Mac, C., 2020. Small Business Owner Race, Liquidity, and Survival. JP Morgan Chase & Co. Institute.

[61] Fershtman, C. and Gneezy, U., 2001. Discrimination in a segmented society: An experimental approach. The Quarterly Journal of Economics, 116(1), pp.351-377.

[62] Fershtman, D. and Pavan, A., 2020. Soft affirmative action and minority recruitment. arXiv preprint arXiv:2004.14953.

[63] Fryer, R.G., Harms, P. and Jackson, M.O., 2018. Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. Journal of the European Economic Association, 17(5), pp.1470-1501.

[64] Fudenberg, D. and Levine, D.K., 1993a. Self-confirming equilibrium. Econometrica, 61(3), pp.523-545.

[65] Fudenberg, D. and Levine, D.K., 1993b. Steady state learning and Nash equilibrium. Econometrica, 61(3), pp.547-573.

[66] Ge, S., Moro, A. and Zhu, B., 2018. Testing for Asymmetric Employer Learning and Statistical Discrimination. SSRN working paper 3243966.

[67] Gentzkow, M. and Shapiro, J.M., 2006. Media bias and reputation. Journal of Political Economy, 114(2), pp.280-316.

[68] Gittins, J. and Jones, D., 1974. A dynamic allocation index for the sequential design of experiments. Progress in statistics, 1(1), pp.241-266.

[69] Giuliano, L., Levine, D.I. and Leonard, J., 2009. Manager race and the race of new hires. Journal of Labor Economics, 27(4), pp.589-631.

[70] Giuliano, L. and Ransom, M.R., 2013. Manager ethnicity and employment segregation. ILR Review, 66(2), pp.346-379.

[71] Glover, D., Pallais, A. and Pariente, W., 2017. Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores. The Quarterly Journal of Economics, 132(3), pp.1219-1260.

[72] Gneezy, U., Imas, A. and List, J., 2015. Estimating individual ambiguity aversion: A simple approach (No. w20982). National Bureau of Economic Research.

[73] Goette, L., Huffman, D., Meier, S. and Sutter, M., 2012. Competition between organizational groups: Its impact on altruistic and antisocial motivations. Management Science, 58(5), pp.948-960.

[74] Goldin, C.D., 1991. The role of World War II in the rise of women's employment. American Economic Review, 81(4), pp.741-756.

[75] Goodman, J.K., Cryder, C.E. and Cheema, A., 2013. Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. Journal of Behavioral Decision Making, 26(3), pp.213-224.

[76] Grether, D.M., 1980. Bayes rule as a descriptive model: The representativeness heuristic. The Quarterly Journal of Economics, 95(3), pp.537-557.

[77] Guarino, C., Reckase, M., Stacy, B. and Wooldridge, J., 2015. A comparison of student growth percentile and value-added models of teacher performance. Statistics and Public Policy, 2(1), pp.1-11.

[78] Guenzel, M. and Malmendier, U., 2020. Behavioral corporate finance: The life cycle of a CEO career (No. w27635). National Bureau of Economic Research.

[79] Guryan, J. and Charles, K.K., 2013. Taste-based or statistical discrimination: the economics of discrimination returns to its roots. The Economic Journal, 123(572), pp.F417-F432.

[80] Headd, B., 2000. The characteristics of small-business employees. Monthly Lab. Rev., 123, p.13.

[81] Hellerstein, J.K., Neumark, D. and Troske, K.R., 2002. Market Forces and Sex Discrimination. Journal of Human Resources, 37(2), pp.353-380.

[82] Hoelzemann, J. and Klein, N., 2018. Bandits in the Lab. Cahiers de recherche 2018-09, Université de Montréal, Département de sciences économiques.

[83] Hoffman, M., Kahn, L.B. and Li, D., 2018. Discretion in hiring. Quarterly Journal of Economics, 133(2), pp.765-800.

[84] Holzer, H.J., 1998. Why do small establishments hire fewer blacks than large ones?. Journal of Human Resources, 33(4), pp.896-914.

[85] Hong, H. and Stein, J.C., 2007. Disagreement and the stock market. Journal of Economic perspectives, 21(2), pp.109-128.

[86] Hossain, T. and Okui, R., 2013. The binarized scoring rule. Review of Economic Studies, 80(3), pp.984-1001.

[87] Kahn, L.B., 2013. Asymmetric information between employers. American Economic Journal: Applied Economics, 5(4), pp.165-205.

[88] Kahn, L.B. and Lange, F., 2014. Employer learning, productivity, and the earnings distribution: Evidence from performance measures. The Review of Economic Studies, 81(4), pp.1575-1613.

[89] Kahneman, D., 2003. Maps of bounded rationality: Psychology for behavioral economics. American Economic Review, 93(5), pp.1449-1475.

[90] Keller, G., Rady, S. and Cripps, M., 2005. Strategic experimentation with exponential bandits. Econometrica, 73(1), pp.39-68.

[91] Kelly, D.J., Quinn, P.C., Slater, A.M., Lee, K., Ge, L. and Pascalis, O., 2007. The other-race effect develops during infancy: Evidence of perceptual narrowing. Psychological Science, 18(12), pp.1084-1089.

[92] Kirschenman, J. and Neckerman, K.M., 1991. “We'd love to hire them, but...”: The meaning of race for employers. In Jencks, C. and Peterson, P.E. (ed.) The urban underclass. Brookings Institution Press, pp.203-232.

[93] Komiyama, J. and Noda, S., 2020. On Statistical Discrimination as a Failure of Social Learning: A Multi-Armed Bandit Approach. arXiv preprint arXiv:2010.01079.

[94] Kurtulus, F.A., 2015. The impact of affirmative action on the employment of minorities and women over three decades: 1973-2003. Upjohn Institute Working Paper No. 15-221.

[95] Landsman, R., 2018. Gender differences in executive departure. Unpublished. Bucknell University.

[96] Lang, K., 1986. A language theory of discrimination. The Quarterly Journal of Economics, 101(2), pp.363-382.

[97] Lang, K., Kahn-Lang Spitzer, A., 2020. Racial discrimination: An economic perspective. Forthcoming in Journal of Economic Perspective.

[98] Lang, K. and Lehmann, J.Y.K., 2012. Racial discrimination in the labor market: Theory and empirics. Journal of Economic Literature, 50(4), pp.959-1006.

[99] Lang, K. and Manove, M., 2011. Education and labor market discrimination. American Economic Review, 101(4), pp.1467-96.

[100] Lange, F., 2007. The speed of employer learning. Journal of Labor Economics, 25(1), pp.1-35.

[101] Laouénan, M. and Rathelot, R., 2017. Ethnic discrimination on an online marketplace of vacation rental. HAL-01514713 Working Paper.

[102] Lazear, E.P., Shaw, K.L. and Stanton, C., 2016. Making do with less: working harder during recessions. Journal of Labor Economics, 34(S1), pp.S333-S360.

[103] Lesner, R.V., 2018. Testing for statistical discrimination based on gender. Labour, 32(2), pp.141-181.

[104] Leung, M.D., 2017. Learning to hire? Hiring as a dynamic experiential learning process in an online market for contract labor. Management Science, 64(12), pp.5651-5668.

[105] Lindqvist, E. and Vestman, R., 2011. The labor market returns to cognitive and noncognitive ability: Evidence from the Swedish enlistment. American Economic Journal: Applied Economics, 3(1), pp.101-28.

[106] Lundberg, S.J. and Startz, R., 1983. Private discrimination and social intervention in competitive labor market. American Economic Review, 73(3), pp.340-347.

[107] March, J.G., 1991. Exploration and exploitation in organizational learning. Organization science, 2(1), pp.71-87.

[108] Meyer, R.J. and Shi, Y., 1995. Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. Management Science, 41(5), pp.817-834.

[109] Miller, C., 2017. The persistent effect of temporary affirmative action. American Economic Journal: Applied Economics, 9(3), pp.152-90.

[110] Mobius, M., Niederle, M., Niehaus, P. and Rosenblat, T., 2014. Managing self-confidence: Theory and experimental evidence (No. w17014). National Bureau of Economic Research.

[111] Mobius, M.M. and Rosenblat, T.S., 2006. Why beauty matters. American Economic Review, 96(1), pp.222-235.

[112] Mobius, M., Rosenblat, T. and Wang, Q., 2016. Ethnic discrimination: Evidence from China. European Economic Review, 90, pp.165-177.

[113] Moore, D.A., Tenney, E.R. and Haran, U., 2015. Overprecision in judgment. The Wiley Blackwell Handbook of Judgment and Decision Making, pp.182-209

[114] Morgan, J. and Várdy, F., 2009. Diversity in the Workplace. American Economic Review, 99(1), pp.472-85.

[115] Moro, A. and Norman, P., 2004. A general equilibrium model of statistical discrimination. Journal of Economic Theory, 114(1), pp.1-30.

[116] Morris, C.N., 1983. Parametric empirical Bayes inference: theory and applications. Journal of the American Statistical Association, 78(381), pp.47-55.

[117] Mortensen, K. and Hughes, T.L., 2018. Comparing Amazon's Mechanical Turk platform to conventional data collection methods in the health and medical research literature. Journal of General Internal Medicine, 33(4), pp.533-538.

[118] Moss, P. and Tilly, C., 2001. Why opportunity isn't knocking: Racial inequality and the demand for labor. Urban Inequality: Evidence from Four Cities. New York, NY: Russel Sage, pp.444-495.

[119] Neumark, D., 2018. Experimental research on labor market discrimination. Journal of Economic Literature, 56(3), pp.799-866.

[120] Nickerson, R.S., 1998. Confirmation bias: A ubiquitous phenomenon in many guises. Review of General Psychology, 2(2), pp.175-220.

[121] Oh, D., Dotsch, R., Porter, J. and Todorov, A., 2020. Gender biases in impressions from faces: Empirical studies and computational models. Journal of Experimental Psychology: General, 149(2), p.323.

[122] Olivola, C.Y. and Todorov, A., 2010. Fooled by first impressions? Reexamining the diagnostic value of appearance-based inferences. Journal of Experimental Social Psychology, 46(2), pp.315-324.

[123] Pager, D. and Karafin, D., 2009. Bayesian bigot? Statistical discrimination, stereotypes, and employer decision making. The Annals of the American Academy of Political and Social Science, 621(1), pp.70-93.

[124] Pager, D., Western, B., Bonikowski, B., Manza, J. and Sauder, M., 2009. Race at Work: The Realities of Race and Criminal Record in the NYC Job Market. Inequality and Society: Social Science Perspectives on Social Stratification. New York: W.W. Norton & Company, pp.425-431.

[125] Paluck, E.L., Green, S.A. and Green, D.P., 2019. The contact hypothesis reevaluated. Behavioural Public Policy, 3(2), pp.129-158.

[126] Pettigrew, T.F. and Tropp, L.R., 2006. A meta-analytic test of intergroup contact theory. Journal of personality and social psychology, 90(5), p.751.

[127] Phelps, E.S., 1972. The statistical theory of racism and sexism. The American Economic Review, 62(4), pp.659-661.

[128] Rabin, M., 2002. Inference by believers in the law of small numbers. The Quarterly Journal of Economics, 117(3), pp.775-816.

[129] Reuben, E., Sapienza, P. and Zingales, L., 2014. How stereotypes impair women's careers in science. Proceedings of the National Academy of Sciences, 111(12), pp.4403-4408.

[130] Robbins, H., 1952. Some aspects of the sequential design of experiments. Bulletin of the American Mathematical Society, 58(5), pp.527-535.

[131] Sarsons, H., 2019. Interpreting signals in the labor market: Evidence from medical referrals. Unpublished. University of Toronto.

[132] Schary, M.A., 1991. The probability of exit. The RAND Journal of Economics, 22(3), pp.339-353.

[133] Scheinkman, J.A. and Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of political Economy, 111(6), pp.1183-1220.

[134] Stephens-Davidowitz, S., 2014. The cost of racial animus on a black candidate: Evidence using Google search data. Journal of Public Economics, 118, pp.26-40.

[135] Sutter, M., 2009. Individual behavior and group membership: Comment. American Economic Review, 99(5), pp.2247-57.

[136] Tajfel, H., Billig, M.G., Bundy, R.P. and Flament, C., 1971. Social categorization and intergroup behaviour. European Journal of Social Psychology, 1(2), pp.149-178.

[137] Tversky, A. and Kahneman, D., 1971. Belief in the law of small numbers. Psychological Bulletin, 76(2), p.105.

[138] U.S. 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]

[139] Van Dalen, H.P. and Henkens, K., 2019. Do stereotypes about older workers change? A panel study on changing attitudes of managers. Forthcoming at International Journal of Manpower.

[140] Waldman, M., 1984. Job assignments, signalling, and efficiency. The RAND Journal of Economics, 15(2), pp.255-267.

[141] Wilson, W.J., 1996. When Work Disappears: The World of the New Urban Poor. New York: Alfred A. Knopf.