Online Appendix for "Experience-based Discrimination"

Louis-Pierre Lepage

Contents

1 Additional Model Discussion

1.1 Certainty about Group A's Productivity

Note that the proofs of Propositions 1-3 do not rely on employers being certain about the productivity of Group A. At one extreme, the results hold directly when allowing for arbitrarily small uncertainty about Group A's productivity. At the other extreme, even if the initial level of uncertainty is identical across groups, if group B is a minority, then most employers will become more uncertain about their productivity over time and the mechanism may operate similarly. Therefore, the extent of discrimination against group B may increase with the degree of relative uncertainty about their productivity, but certainty about group A's productivity is not necessary to generate the model's predictions.

1.2 Signals of Individual Productivity and Endogenous Worker Investments

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. Negatively-biased beliefs about the mean productivity of group B conditional on a given signal value arise as in the baseline model. Since employers above the hiring cutoff are willing to pay more for a group B worker conditional on s_i, 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).

When groups are ex-ante equally productive, statistical discrimination models usually generate outcome disparities because 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.

While a formal model of endogenous worker investment is beyond the scope of this paper, in the 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. Group B may still be incentivized to sort into occupations where the information asymmetry problem faced by employers is lesser, providing a rationale for group specialization. Similarly, if group B earns lower expected returns from the labor market overall, they may have incentives to invest less in human capital, which could exacerbate discrimination.

1.3 Firm Size and Hiring Policy

Larger employers who hire more have a higher value of learning and should learn more quickly. Negative biases may be less likely to persist, and these employers would be predicted to hire a higher fraction of group B workers, consistent with evidence reported in Miller (2017) for black workers. These implications relate to large establishments with centralized human resources (HR) services rather than large firms with decentralized hiring. When the hiring process is decentralized, individual managers have been shown to play an important role in the group composition of hires (Giuliano et al., 2009; Benson and Lepage, 2022) and common policies like pre-employment testing or hiring algorithms typically fail to address concerns of endogenous learning specifically (Bergman et al., 2020).

Implications for the model predictions remain limited if each establishment hires a negligible fraction of the labor force and there is size heterogeneity above the hiring cutoff. Unless all of group B is hired by large establishments with centralized hiring, then these establishments are not marginal, by definition, and the wage is determined by smaller establishments who learn more slowly. 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.

2 Model Simulations and Comparative Dynamics

To illustrate the model's dynamics, a set of simulations was computed over 1,000 periods with 10,000 employers and 10,000 workers. I consider a relative size for group B of 25%. Given a prior distribution of beliefs, the initial market-clearing wage where employers maximize their expected profits is found. Beliefs are updated such that those above 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.

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, which do not substantively vary with parameter choice.[1]

Panel A of Figure A2-1 shows the evolution of beliefs for key moments of the distribution, without entry and exit. 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 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 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.

Figure A2-2 presents simulations with market entry and exit to illustrate Remark 1. I set the firm exit rate weighted by the share of employment at 2% per year (Crane et al., 2022). A standard estimate for the labor cost share is around 0.6, which combined with a group B share of 0.25, yields an exit rate differential of 15% for employers below versus above the hiring cutoff for group B. 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 with market exit is that, since all employers exit the market in finite time, some employers above the hiring cutoff always have negatively-biased beliefs.

To show how the wage gap varies with exit rates and differential exit rates, I show simulations comparing aggregate exit rates of 2% and 1% in Panel A of Figure A2-3 and simulations comparing exit rate differentials of 15, and 100% in Panel B. These simulations indicate that the wage gap decreases with higher market exit rates as well as higher differential market exit rates for employers below the hiring cutoff for group B.

The expected size of the wage gap is influenced by other parameters as displayed in Figure A2-4. A larger relative size for group B leads to a lower relative wage for the group. A lower mean productivity for group B also leads to a lower wage. Negatively-biased priors initially decrease the group B wage, but have little impact in the long run. A higher employer prior precision or lower variance in productivity of group B increase the wage. Assuming homogeneous rather than unbiased employer priors has little impact on the wage (slightly higher), while introducing stereotype bias through employers overestimating their signal

[1] Similarly, the initial state 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.

precision (or equivalently underestimating the variance in group B's productivity) decreases the wage.

Similarities and differences between the simulated wage path and empirical wage trends naturally do not provide a test of the model. 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, Figure A2-4 shows that negatively-biased priors can generate a group B wage which starts and remains below that of group A, but increases over time. Similarly, in the baseline model, employers begin by hiring group B most often and gradually decrease their hiring of the group, but the simulation with negatively-biased priors predicts the opposite pattern, plausibly more in line with historical trends. More generally, the simulations should be interpreted as a way to visualize model dynamics, rather than attempt to quantify the extent of discrimination in practice.

Figure A2-1: Model Simulation without Entry and Exit

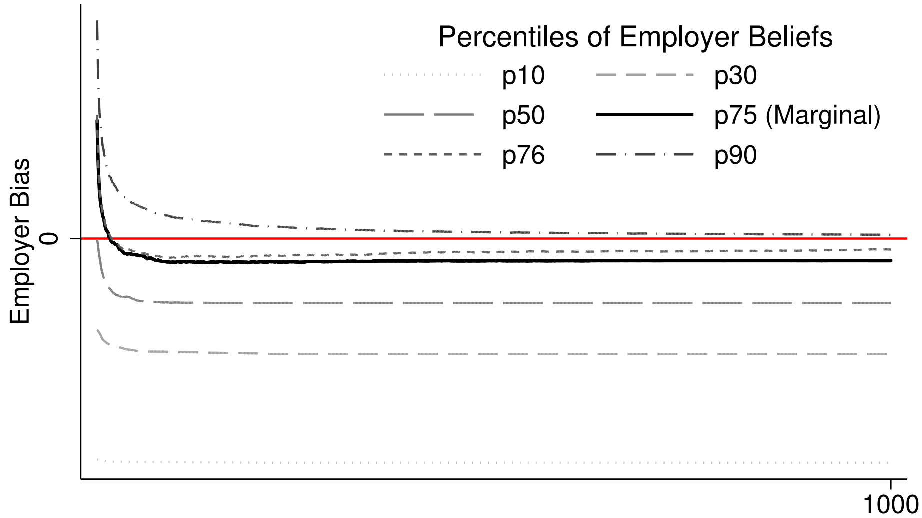

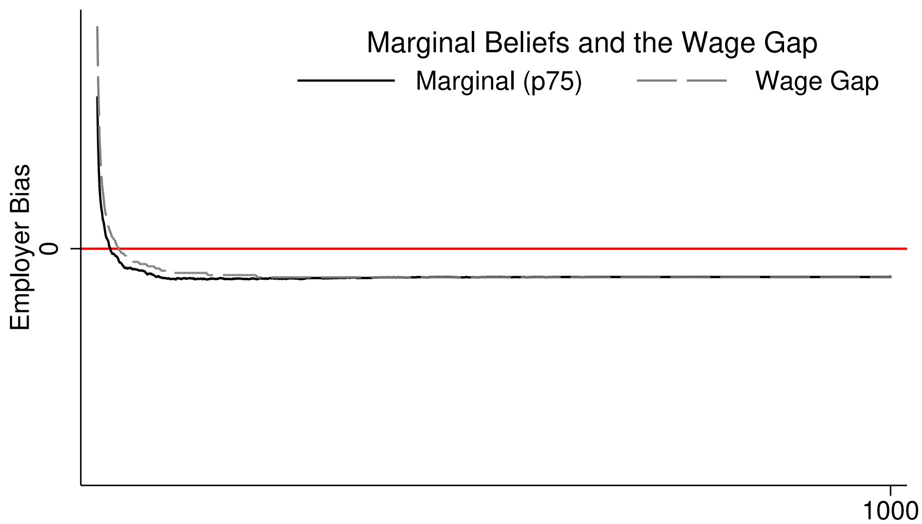

The fraction of group B workers is 0.25. Worker productivity is distributed $N(0,2)$, prior beliefs are distributed $N(0,1)$. w_A is normalized to 0 and β is set to 0.9.

Figure A2-2: Model Simulation with Market Entry and Exit, 15% Exit Differential

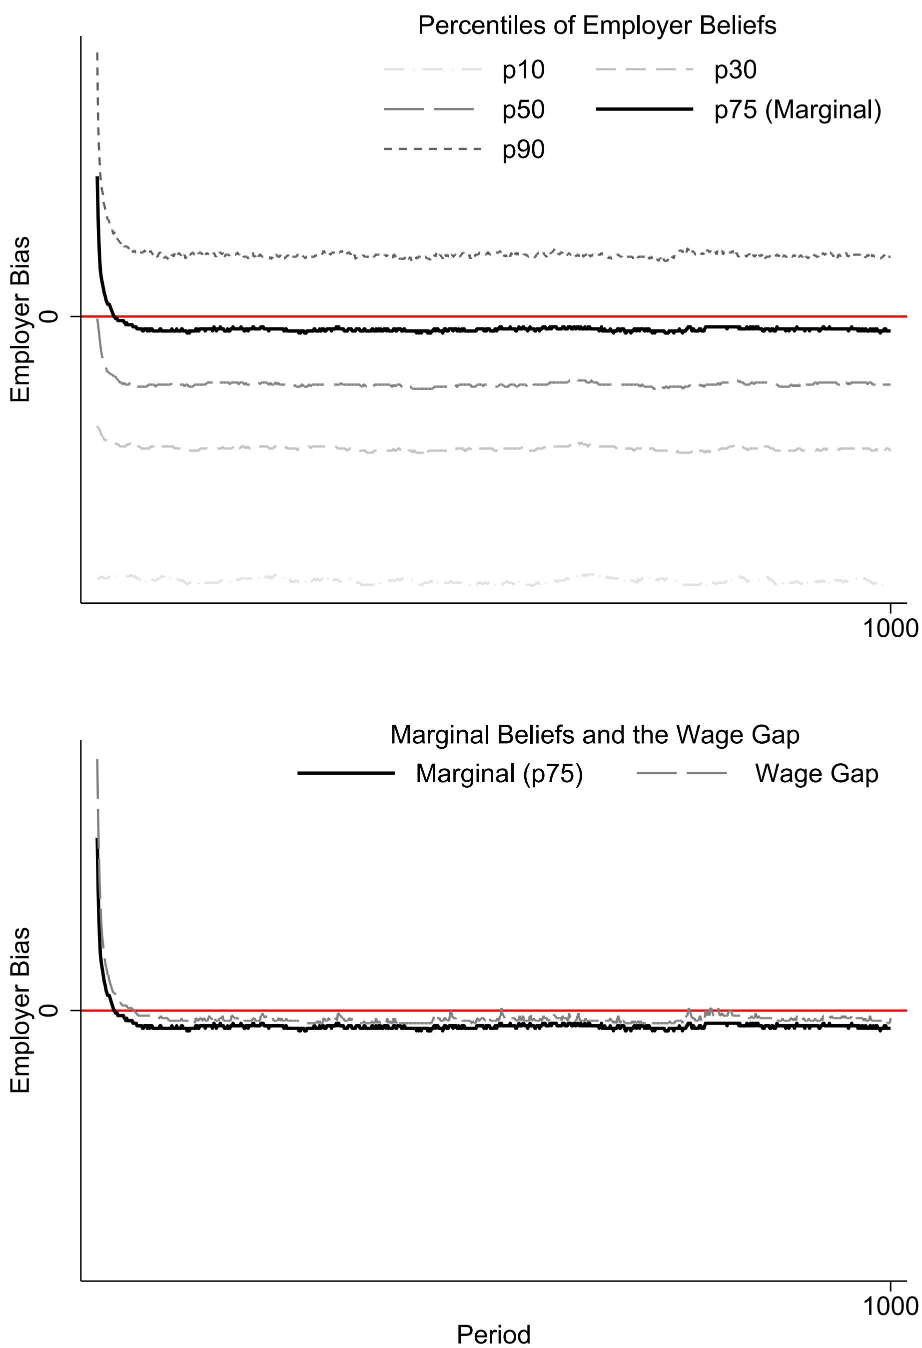

The aggregate exit rate corresponds to 2% each period, with a 15% higher exit rate for employers below the hiring cutoff for group B. New entrants have mean beliefs equal to 0 (unbiased). See Figure A2-1 for other parameter choices.

Figure A2-3: Wage Gap and Competition

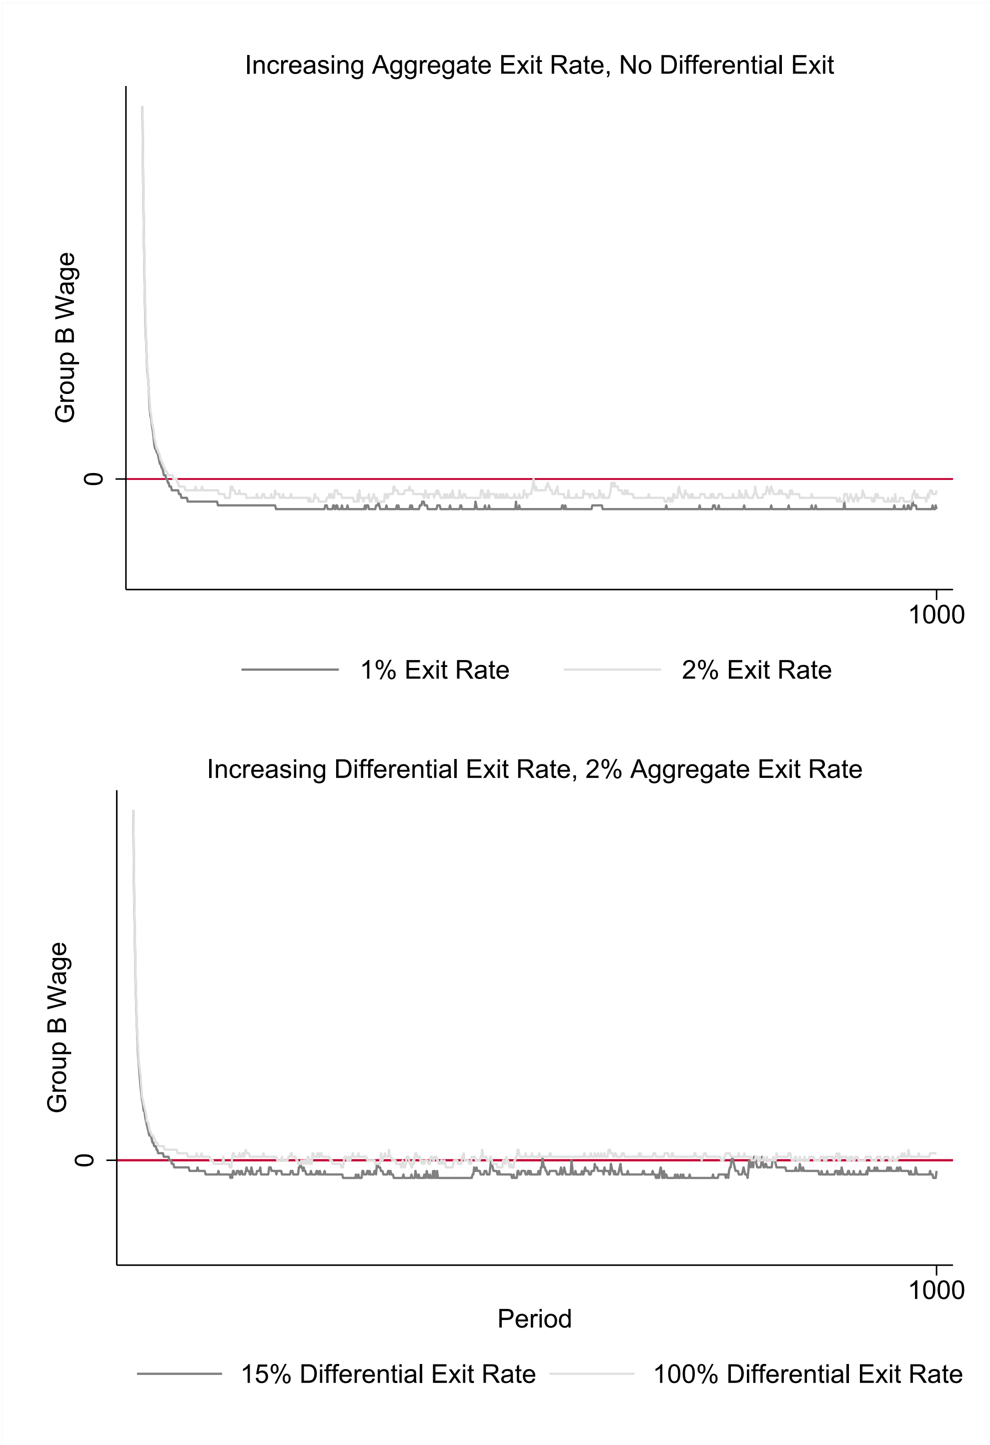

The aggregate exit rates correspond to 1% and 2% each period for Panel A, with no differential exit rate for employers below the hiring cutoff for group B. The aggregate exit rate corresponds to 2% each period for Panel B, with differential exit rates of 15% and 100% for employers below the hiring cutoff for group B. New entrants have mean beliefs equal to 0 (unbiased). See Figure A2-1 for other parameter choices.

Figure A2-4: Wage Gap and Model Parameters

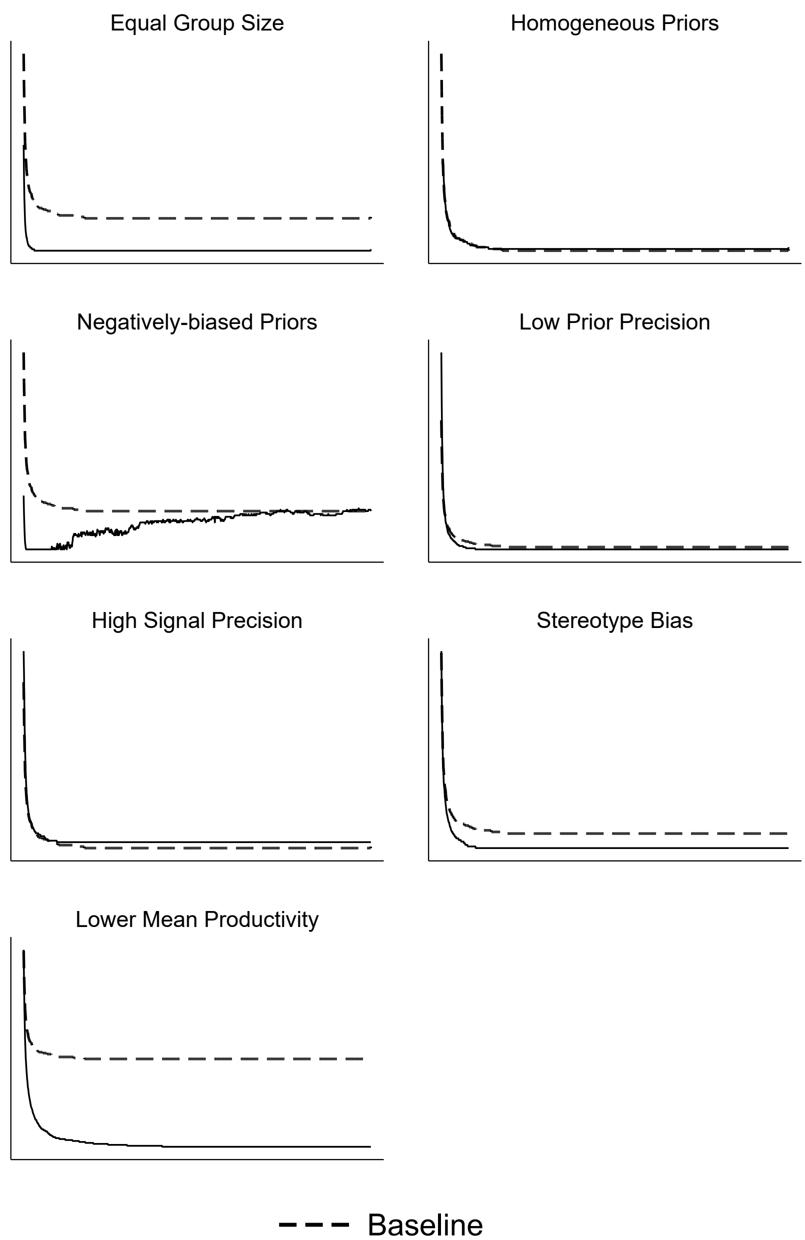

Equal Group Size refers to group B being of equal size to group A (50% of workers). Homogeneous Priors refers to each employer holding prior $\mu_0 = 0$. Negatively-Biased Priors refers to employers having mean prior beliefs below the true value (-1 vs 0). Low Prior Precision corresponds to a case with 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). See Figure A2-1 for other parameter choices.

3 Additional Experiment Information and Results

3.1 Recruitment and Implementation

An exchange rate of 1,000 credits for \$0.2 was used. The subject pool was restricted to US adults with an approval rating of above 95% and at least 100 completed tasks. Employers also had to answer comprehension questions to ensure a good understanding of every aspect of the task.[2] The experiment was implemented using oTree (Chen et al., 2016).

Workers earned 250 credits per puzzle solved. They 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 \$10-\$12.

Employers earned 220 credits per puzzle solved by their worker each period, paid for a random subset H of 5 periods. Belief elicitation was made operational as follows. Employers reported their beliefs μ_{Bjt} about the group's mean productivity. Then, each period, beliefs were used to compute a squared prediction error $(\mu - \mu_{Gjt})^2$. A set of two periods R was randomly selected for payment. If the period was selected for payment, employers received 110 credits if their squared prediction error was below or equal to 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 B productivity, but the belief elicitation procedure isolates learning about the mean to focus on the impact of experiences on mean posterior beliefs. Similarly, employers were not given information on the minimum and maximum number of puzzles solved by workers to keep the instructions as simple and brief as possible and because including or omitting this information does not alter the framework's theoretical predictions. The total payoff of employer j corresponds to

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

where y_{it} is their period t hire's productivity. 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 \$12-\$15.[3] Based on power calculations and pilot experiments, 297 employers were assigned to Treatment *Baseline*, 135 to Treatment *Control*, 148 to Treatment *Exploration*, 152 to Treatment *Equal*, 138 to Treatment *Information*, 239 to Treatment *Gender*, and 190 to Treatment

[2]They could attempt to answer the questions as many times as they wished within a one hour period, but could not continue without answering all questions correctly. Over 60% of participants did not complete the questions and abandoned the experiment, substantially improving data quality. Other tests of quality included investigating IP address clustering and string-based attention questions.

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

Elicitation which had their beliefs elicited only at the end of the hiring task.[4] Balance tests across treatments are presented in Table A3-2.[5] See Table A3-1 for a summary of employer treatments.

3.2 Example Worker Puzzle

Figure A3-1: Example Puzzle

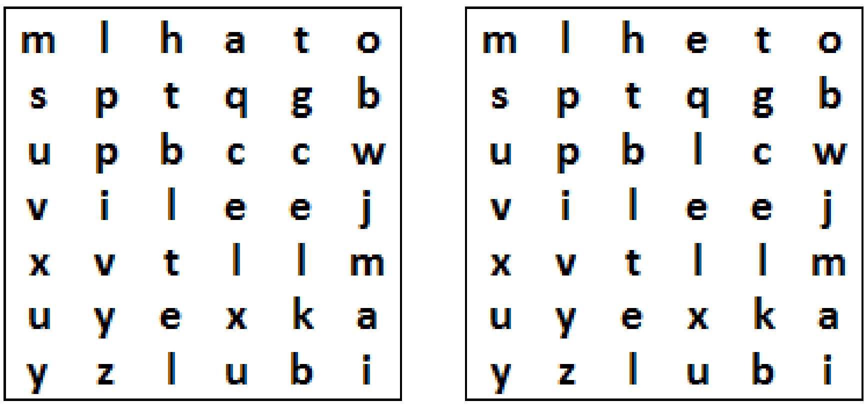

Note. The square with characters on the right differs from the square on the left in two letters. Workers had to identify those letters to solve the puzzle.

[4] These numbers exclude employers who reported beliefs above (below) the minimum number of puzzles solved by workers or failed other basic data quality checks, namely not updating beliefs, systematically updating in the wrong direction, or not updating as a function of their productivity draws. These exclusions ensure that the results are not driven by outlier unrealistic beliefs.

[5] MTurk sessions corresponding to different employer treatments were conducted at different times, but Table A3-2 shows little difference in characteristics across treatments and Table A3-5 shows little difference in behavior across employer characteristics within the *Baseline* treatment.

3.3 Difference in Hiring and Final Beliefs

Figure A3-2: Difference in Final Employer Beliefs, *Baseline* versus *Control* Treatments

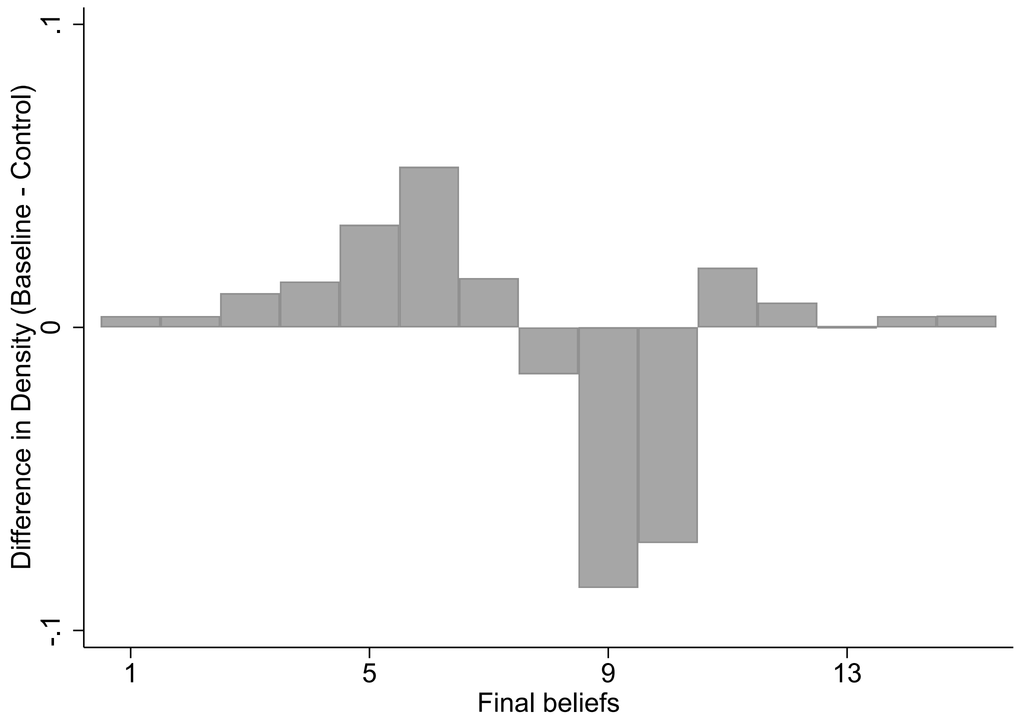

Note. The figure plots the difference in the final belief distribution between the *Baseline* and *Control* treatments. See Figures 1-3 for additional details.

Figure A3-3: Difference in Final Employer Beliefs by Total B Hires and Period of Last B Hire, *Baseline* Treatment

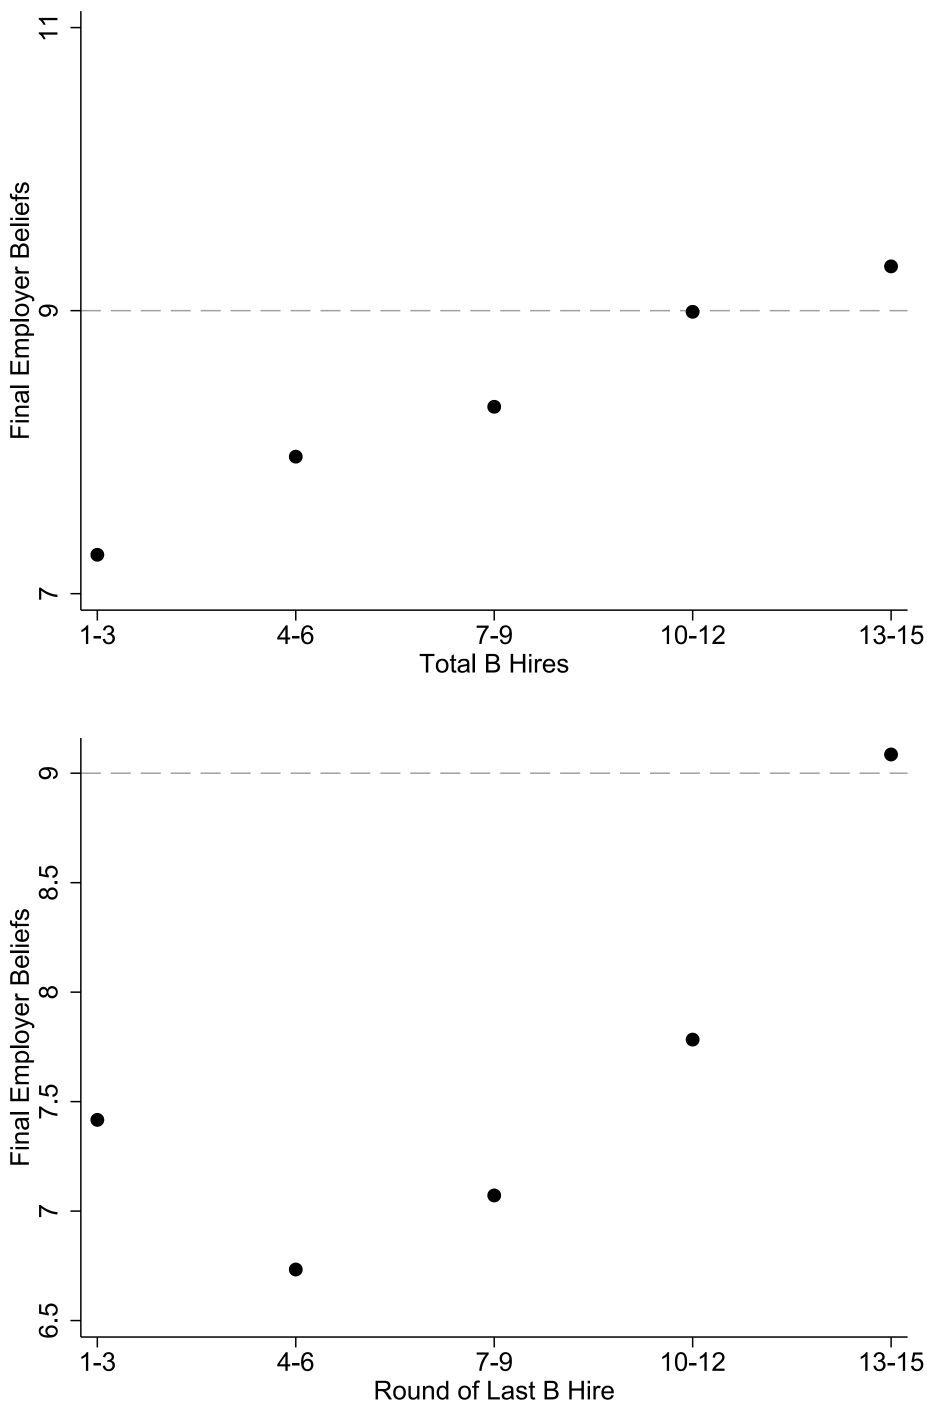

Note. See Figure 1 for additional details.

Figure A3-4: Impact of First Negative Experience with Group B on Hiring and Beliefs, by Productivity of the First B Hire, *Baseline* Treatment

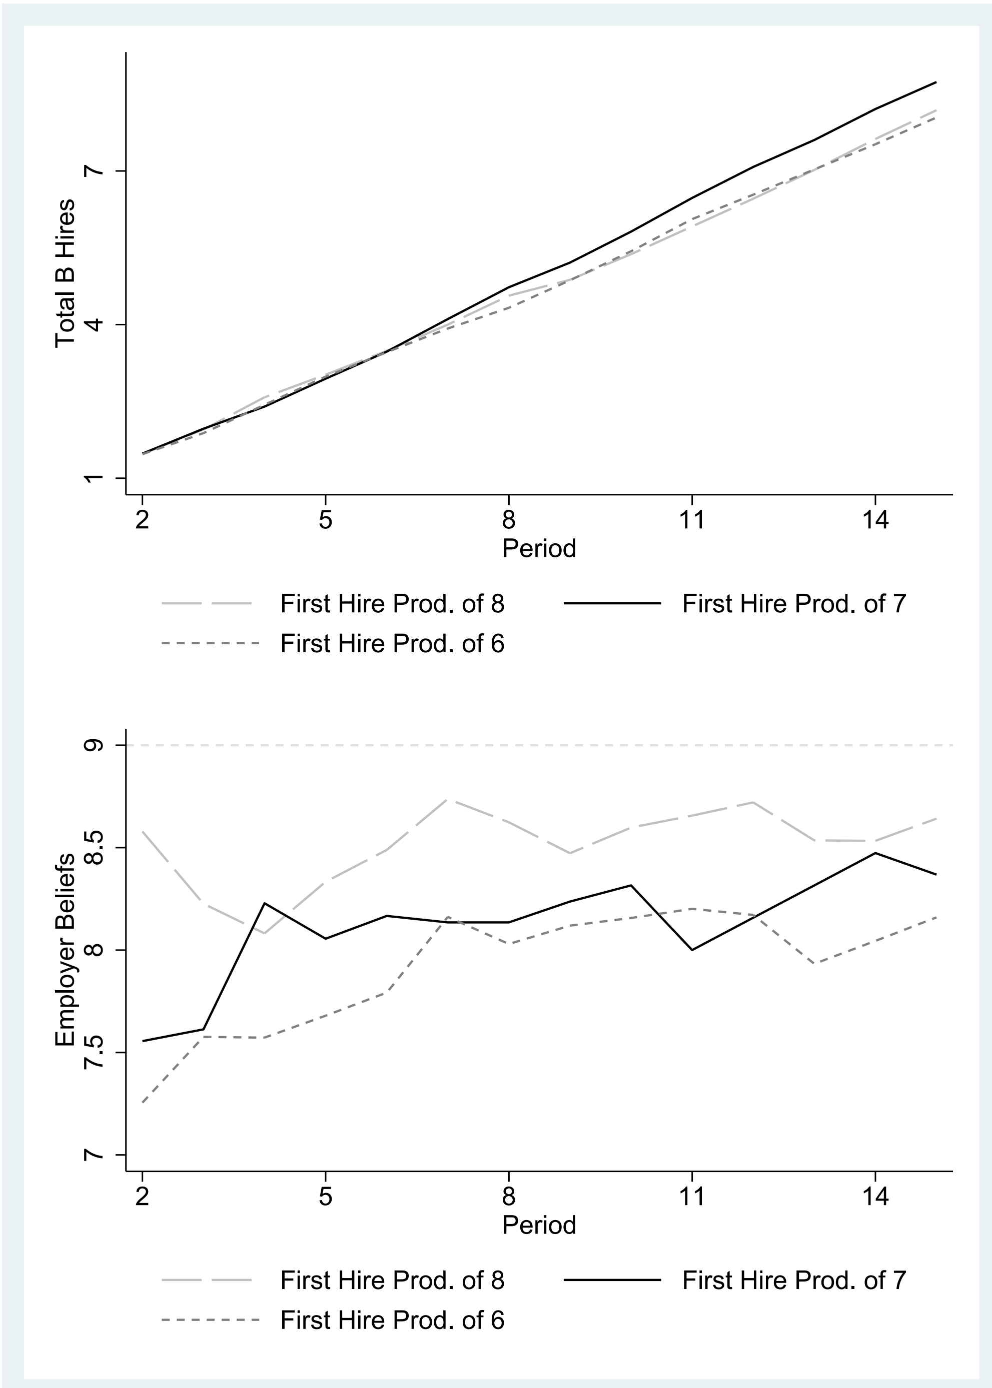

Note. See Figure 2 for additional details.

3.4 Summary of Employer Treatments

Table A3-1: Employer Treatments

	N	Hiring	Minority Status	Belief Elicitation	Additional information
Baseline	297	Group B or A	Group B Minority	Prior, B hire	N/A
Control	135	Group B	N/A	Every period	N/A
Exploration	148	Group B or A	Group B Minority	Prior, B hire	Extra credits for hiring B
Equal	152	Group B or A	Equal Group Sizes	Prior, B hire	N/A
Information	138	Group B or A	Group B Minority	Prior, B hire, period 10-15	Group B, periods 10-15
Gender	239	Female or Male	Female Minority	Prior, female hire	N/A
Elicitation	190	Group B or A	Group B Minority	End	N/A

Note. Group B has 50 workers. When both groups are presented as equally-sized, group A also has 50 workers. Otherwise, group A has 150 workers. Female workers represent 77 out of 200 workers. The additional information given to Treatment *Information* corresponds to the average productivity of 5 randomly-selected group B workers previously hired by other employers for each period from 10 to 15.

3.5 Balance Tests

Table A3-2: Employer Characteristics Across Treatments

	Treatment	Mean	SD	N	Difference with baseline	Joint difference		Treatment	Mean	SD	N	Difference with baseline	Joint difference
	Baseline	35.29	10.31	297				Baseline	0.09	0.28	297		
	Control	35.21	9.53	135	0.94			Control	0.08	0.28	135	0.84	
	Exploration	36.89	11.04	148	0.13			Exploration	0.07	0.26	148	0.64	
Age	Elicitation	36.72	10.50	190	0.14	0.41	Asian	Elicitation	0.09	0.29	190	0.94	0.66
	Equal	35.61	11.32	152	0.76			Equal	0.05	0.22	152	0.19	
	Information	36.76	9.86	138	0.16			Information	0.07	0.25	138	0.43	
	Gender	36.78	10.45	239	0.100			Gender	0.05	0.23	239	0.14	
	Baseline	0.66	0.48	297				Baseline	0.07	0.26	297		
	Control	0.59	0.49	135	0.20			Control	0.08	0.28	135	0.79	
	Exploration	0.65	0.48	148	0.87			Exploration	0.04	0.20	148	0.17	
Male	Elicitation	0.60	0.49	190	0.17	0.78	Hispanic	Elicitation	0.05	0.21	190	0.24	0.62
	Equal	0.61	0.49	152	0.29			Equal	0.05	0.21	152	0.25	
	Information	0.62	0.49	138	0.50			Information	0.06	0.24	138	0.54	
	Gender	0.63	0.48	239	0.55			Gender	0.05	0.23	239	0.36	
	Baseline	0.73	0.45	297				Baseline	0.85	0.36	297		
	Control	0.72	0.45	135	0.85			Control	0.84	0.36	135	0.99	
	Exploration	0.74	0.44	148	0.72			Exploration	0.83	0.38	148	0.70	
White	Elicitation	0.75	0.43	190	0.54	0.51	College	Elicitation	0.82	0.39	190	0.40	0.69
	Equal	0.78	0.41	152	0.20			Equal	0.87	0.34	152	0.51	
	Information	0.73	0.45	138	0.92			Information	0.88	0.33	138	0.38	
	Gender	0.80	0.41	239	0.07			Gender	0.87	0.34	239	0.49	
	Baseline	0.09	0.29	297				Baseline	0.72	0.45	297		
	Control	0.07	0.26	135	0.49			Control	0.72	0.45	135	0.98	
	Exploration	0.07	0.26	148	0.48			Exploration	0.68	0.47	148	0.45	
Black	Elicitation	0.10	0.30	190	0.84	0.92	Employment	Elicitation	0.74	0.44	190	0.55	0.92
	Equal	0.09	0.28	152	0.76			Equal	0.72	0.45	152	0.89	
	Information	0.10	0.30	138	0.81			Information	0.70	0.46	138	0.76	
	Gender	0.08	0.26	239	0.44			Gender	0.74	0.44	239	0.62	

Note. "Difference with Baseline Treatment" presents p-values from pairwise t-tests of equal sample means between the Baseline treatment and other treatments. "Joint Difference" presents p-values from multiple-comparison tests using one-way analysis-of-variance models. See Table A3-1 for a description of treatments.

3.6 Additional Evidence for Exploration Treatment

Table A3-3: Differences in Hiring between the *Baseline* and the *Exploration* Treatments

	Subsequent B Hiring		
	(1)	(2)	(3)
Prod. of First Hire	0.105		
	(0.021)		
Exploration * Prod. of First Hire	-0.135		
	(0.035)		
Pos. Exp. with First Hire		0.727	
		(0.153)	
Exploration * Pos. Exp. with First Hire		-0.985	
		(0.262)	
Neg. Exp. with First Hire			-0.617
			(0.154)
Exploration * Neg. Exp. with First Hire			0.805
			(0.266)
Outcome Mean	5.329	5.329	5.329
N. Obs.	3,947	3,947	3,947

Note. Robust standard errors are presented in parentheses. Treatment *Exploration*: as in Treatment *Baseline*, but employers are given a 440 credit bonus each period they hire from group B. See Tables 2 and 3 for additional details.

3.7 Additional Evidence for Equal Treatment

Table A3-4: Differential Impact of B Hires on Final Bias, *Baseline* versus *Equal* Treatments

	Final Bias		
	(1)	(2)	(3)
Prod. of First Hire	-0.087		
	(0.038)		
Equal * Prod. of First Hire	0.135		
	(0.063)		
Pos. Exp. with First Hire		-0.334	
		(0.192)	
Equal * Pos. Exp. with First Hire		0.574	
		(0.333)	
Neg. Exp. with First Hire			0.682
			(0.190)
Equal * Neg. Exp. with First Hire			-0.615
			(0.335)
Outcome Mean	1.92	1.92	1.92
N. Obs.	403	403	403

Note. Robust standard errors presented in parentheses. See Tables 2 and 3 for additional details.

3.8 Heterogeneity Across Employer Characteristics

Table A3-5: Differences in Hiring and Bias by Employer Characteristic, Treatment *Baseline*

	Total B Hires (1)	Final Bias (2)
Prejudice	-1.493	0.068
	(0.451)	(0.115)
High School	-0.613	-0.222
	(0.832)	(0.220)
Age	0.024	0.007
	(0.028)	(0.009)
Male	0.285	-0.109
	(0.619)	(0.196)
Employed	-0.037	-0.102
	(0.652)	(0.207)
Black	-1.233	0.797
	(0.919)	(0.354)
Hispanic	-0.025	0.097
	(1.166)	(0.333)
Outcome Mean	8.04	1.67
N. Obs.	297	297

Note. 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. 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. Employed is an indicator variable for whether the participant is employed beyond their work on Mechanical Turk. See Tables 2 and 3 for additional details.

3.9 Deviations from Bayesian Updating

Bias formation could be affected by stereotype formation, among other factors. Variance in group B productivity is unknown to employers, but updating about the mean can still be used to infer deviations from Bayesian updating. For every round in which an employer reports their beliefs, I calculate their implied $t = 0$ parameter κ_0, which represents initial beliefs about variance in productivity.[6] Under Bayesian updating, κ_0 is a positive time-invariant constant, with a lower value implying more updating conditional on a signal.

A decreasing κ_0 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 undefined 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, if an employer with prior 9 observes signals of mean 8 and reports posterior beliefs 7. Alternatively, a negative κ_0 can be consistent with over-weighting of positive or negative experiences, such that prior beliefs are closer to $\bar{x}$ than posterior beliefs. For example, if an employer with prior 9 observes signals of mean 8, but reports posterior beliefs 10.

Table A3-6 summarizes implied values of κ_0 and whether they change with experience hiring B or the productivity of the last B hire. Column 1 indicates that κ_0 decreases with hiring experience, consistent with over-updating. It also suggests a small increase in κ_0 and therefore decrease in the extent of updating if the last B hire was more productive. Columns 2-3 indicate that κ_0 is more likely to be negative with experience, primarily reflecting over-updating rather than over-weighting. Around 26% of values are missing, arising from employers often reporting their beliefs as integers. Overall, the results are consistent with employers updating their beliefs by more than a Bayesian benchmark, which amplifies bias formation in theory (see Appendix 2).

[6] 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 $\mu_n = \frac{\kappa_0 \mu_0 + n\bar{x}}{\kappa_0 + n}$. I can recover κ_0 given that everything else is observed.

Table A3-6: Departures from Bayesian Updating, Treatment *Baseline*

	κ_0	Over-Updating	Over-Weighting	Prob. κ_0 Missing
	(1)	(2)	(3)	(4)
Number of Hires	-0.433	0.035	0.010	0.002
	(0.086)	(0.004)	(0.003)	(0.004)
Prev. Hire Prod.	0.067	0.002	-0.003	-0.003
	(0.042)	(0.004)	(0.003)	(0.003)
Mean	-1.28	0.39	0.10	0.25
N. Obs.	1,791	1,791	1,791	2,389

Note. Clustered standard errors at the employer level are presented in parentheses. Regressions include employer fixed effects to capture time-invariant tendencies across employers to hire from a group and update their belief. κ_0 represents $t = 0$ employer beliefs about the variance in productivity of group B recovered from their posterior mean updating. A larger value implies less updating from experiences, and a decreasing value with experience is consistent with employers updating more than implied by their prior about productivity variance. Over-updating corresponds to employers updating too far away from their prior in the direction of the mean signal they observe. Over-weighting of positive or negative experiences corresponds to employers having prior beliefs that are closer to the mean signal they observe than their posterior beliefs. See Table 2 for additional details.

3.10 Impact of the Productivity of Group B Hires on Hiring and Beliefs, No Controls

Table A3-7: Impact of the Productivity of Group B Hires on Hiring and Beliefs

	Subsequent number of group B hires						Final beliefs about group B productivity					
Panel A) *Baseline* treatment	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
Prod. of B hire	0.105						0.089					
	(0.021)						(0.011)					
Prod. of B hire X # of prev. B hires		-0.019						-0.005				
		(0.004)						(0.002)				
Positive Experience			0.727						0.556			
			(0.153)						(0.078)			
Positive Exp. X # of prev. B hires				-0.128						-0.042		
				(0.030)						(0.018)		
Negative Experience					-0.617						-0.555	
					(0.154)						(0.077)	
Negative Exp. X # of prev. B hires						0.140						0.040
						(0.029)						(0.018)
Outcome mean	5.115	5.115	5.115	5.115	5.115	5.115	8.975	8.975	8.975	8.975	8.975	8.975
Number of observations	2,389	2,389	2,389	2,389	2,389	2,389	2,389	2,389	2,389	2,389	2,389	2,389
Panel B) Differential impact, *Baseline* versus *Control*												
Baseline X Prod. of B hire							0.070					
							(0.014)					
Base. X Prod. of B hire X # of prev. B hires								-0.007				
								(0.003)				
Baseline X Positive Experience									0.442			
									(0.096)			
Base. X Positive Exp. X # of prev. B hires										-0.050		
										(0.023)		
Baseline X Negative Experience											-0.410	
											(0.096)	
Base. X Negative Exp. X # of prev. B hires												0.052
												(0.023)
Outcome mean							8.913	8.913	8.913	8.913	8.913	8.913
Number of observations							4,414	4,414	4,414	4,414	4,414	4,414

Note. Robust standard errors are presented in parentheses. Treatment *Baseline*: each period, employers choose between hiring from group A or B. Treatment *Control*: as in Treatment *Baseline*, but employers can only hire from group B each period. Group A is the majority with 75% of workers. Beliefs about the mean productivity of group B are elicited before the first hire and after every hire from the group. Regressions in Panels A and B include an individual measure of ambiguity aversion calculated as in Gneezy et al. (2015) and the employer's prior beliefs about group B's average productivity elicited before the hiring task. A positive (negative) experience refers to a group B hire having productivity above (below) the mean productivity of group A, 9. See Table 2 for additional details.

3.11 Ambiguity Aversion and Hiring

Table A3-8: Impact of Ambiguity Aversion on Hiring and Interaction with First Hire Productivity, Treatment *Baseline*

	Total B Hires (1)	Total B Hires (2)	Total B Hires (3)	Total of 1 B hire (4)	Total of 1 B hire (5)	Total of 1 B hire (6)	Total of 2 B hires (7)	Total of 2 B hires (8)	Total of 2 B hires (9)
Ambiguity Aversion	-0.044	-0.093	-0.004	-0.000	-0.002	-0.000	0.002	0.003	-0.000
	(0.052)	(0.065)	(0.063)	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)
Amb. * Neg. Exp. with First Hire		0.112			0.003			-0.004	
		(0.095)			(0.005)			(0.004)	
Amb. * Pos. Exp. with First Hire			-0.100			-0.001			0.004
			(0.097)			(0.005)			(0.004)
Outcome Mean	8.044	8.981	8.981	0.044	0.049	0.049	0.037	0.041	0.041
N. Obs.	297	266	266	297	266	266	297	266	266

	Total of 3 B hires (10)	Total of 3 B hires (11)	Total of 3 B hires (12)	Total of 4 B hires (13)	Total of 4 B hires (14)	Total of 4 B hires (15)	Total of 5 B hires (16)	Total of 5 B hires (17)	Total of 5 B hires (18)
Ambiguity Aversion	0.003	0.006	0.001	-0.002	-0.004	0.000	0.001	0.001	0.001
	(0.002)	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.002)	(0.003)	(0.003)
Amb. * Neg. Exp. with First Hire		-0.006			0.004			0.001	
		(0.005)			(0.006)			(0.005)	
Amb. * Pos. Exp. with First Hire			0.007			-0.006			-0.001
			(0.005)			(0.006)			(0.005)
Outcome Mean	0.054	0.060	0.060	0.064	0.071	0.071	0.047	0.053	0.053
N. Obs.	297	266	266	297	266	266	297	266	266

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

References

[1] 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.

[2] Benson, A. and Lepage, L., 2022. Learning to Discriminate on the Job. Unpublished. University of Minnesota.

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

[4] 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.

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

[6] Crane, L.D., Decker, R.A., Flaaen, A., Hamins-Puertolas, A. and Kurz, C., 2022. Business exit during the COVID-19 pandemic: Non-traditional measures in historical context. Journal of Macroeconomics, 72, p.103419.

[7] 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.

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

[9] 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.

[10] Lepage, L., 2023. "Experienced-based Discrimination: Dataset". Inter-University Consortium for Political and Social Research, Ann Arbor, MI. AEA Project openICPSR-192292.

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

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