THE ECONOMIC JOURNAL

The Economic Journal, **134** (*February*), 485–514 https://doi.org/10.1093/ej/uead079
Advance Access Publication Date: 27 September 2023

THE BOSS IS WATCHING: HOW MONITORING DECISIONS HURT BLACK WORKERS*

Costas Cavounidis, Kevin Lang and Russell Weinstein

African Americans face shorter employment durations than similar Whites. We hypothesise that employers discriminate in acquiring or acting on ability-relevant information. In our model, monitoring Black, but not White, workers is self-sustaining. New Black hires were more likely fired by previous employers after monitoring. This reduces firms' beliefs about ability, incentivising discriminatory monitoring. We confirm our predictions that layoffs are initially higher for Black than non-Black workers, but that they converge with seniority and decline more with the Armed Forces Qualification Test for Black workers. Two additional predictions, lower lifetime incomes and longer unemployment durations for Black workers, have known empirical support.

> [T]he Black screwup ... faces the abyss after one error, while the White screwup is handed second chances ....
>
> Ibram X. Kendi, *How to Be an Antiracist* (2019, p.93)

Many Americans, especially African Americans, believe that Black workers 'don't get second chances'[1] or face additional scrutiny in the workplace. Similarly, Black workers are admonished to be 'twice as good'[2] in order to succeed. If Black workers are subject to higher standards or scrutinised more heavily, we expect this to be reflected in more separations.

Indeed, the data support the idea of shorter employment duration for Black workers.[3] Bowlus *et al*. (2001) detected and pondered the disparity in job destruction rates; Bowlus and Eckstein (2002) estimated that young Black male high school graduates had roughly 2/3 the job spell duration of their White counterparts.[4] In addition, more of their job spells end in unemployment,

* Corresponding author: Costas Cavounidis, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. Email: c.cavounidis@warwick.ac.uk

This paper was received on 16 August 2021 and accepted on 21 September 2023. The Editor was Barbara Petrongolo.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.8346896.

This research was funded in part by (U.S.) National Science Foundation grants SES-1260917 and SES-1851636. It draws on similar intuition to Cavounidis and Lang (2015), but greatly simplifies the theory and adds an empirical section testing the main new predictions. We are grateful to Cheonghum Park for excellent research assistance and to Gautam Bose, Sambuddha Ghosh, Larry Katz, Ed Lazear, Calvin Luscombe, Ioana Marinescu, Derek Neal, Marina Soley-Bori, Bill Spriggs and participants in seminars and conferences at Boston University, Cornell, CUHK-Shenzhen, EALE, IZA/OECD, NBER, the Royal Economic Society, Tufts, the University of California, Riverside, SOLE, the University of Oklahoma, UQAM and West Point Military Academy for helpful comments and suggestions on this and the earlier paper. The usual caveat applies.

[1] This assertion can be found in a range of occupations, including football coaching (Reid, 2015), music and films (*The Guardian*, 2014) and more generally (Spencer, 2014).

[2] See Coates (2012) and Mabry (2007).

[3] Throughout this paper, we distinguish between employment duration by which we mean the length of an employment spell and job duration by which we mean the time a worker spends with a particular employer. Job duration depends on, among other factors, the arrival rate of outside offers. Our model abstracts from wage renegotiation, but can be modified to incorporate it, as shown in Section 3.5.1.

[4] Using the NLSY data for 1985 and 1988.

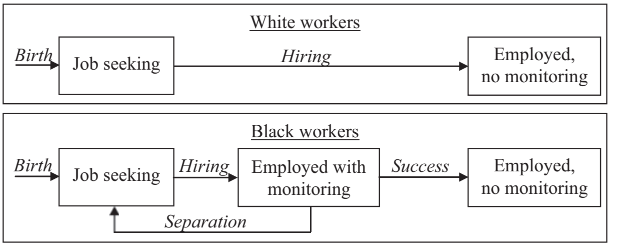

Fig. 1. *White Workers' Perpetual Employment, Black Workers' Churning Cycle.*

suggesting that Black workers have much shorter employment spells. Both papers assume an exogenously higher separation rate for Black workers to fit their models to the data. Lang and Lehmann (2012) showed that differences in unemployment duration alone are insufficient to account for the Black/White unemployment rate gap and, therefore, that Black workers' employment stints are shorter. This aspect of labour discrimination has thus far eluded theoretical explication.

In this paper, our proposed explanation for differential employment durations is, in its broadest sense and consistent with the observations above, that firms discriminate in the acquisition or use of productivity-relevant information. That is, firms either learn differently about Black workers or, when information regarding ability is received, they condition how they act on it on workers' race. Crucially, we establish that such discrimination can be self-perpetuating.

The essence of our model is that, because firms scrutinise Black workers more closely, a larger share of low-performance workers will separate into unemployment. As a result, since productivity is correlated across jobs, the Black unemployment pool is 'churned' and therefore weaker than the White unemployment pool. Since workers can, at least to some extent, hide their employment histories, race serves as an indicator of expected worker productivity. This, in turn, makes monitoring newly hired Black (but not White) workers optimal for firms. Figure 1 illustrates employment in the two labour markets. Our model shares the churning mechanism with Masters (2014), where information acquisition takes the form of exogenous pre-employment signals rather than endogenous monitoring on the job. Bardhi *et al.* (2019) explored endogenous discriminatory employee monitoring for quality, finding that only bad-news monitoring leads to persistent discrimination when worker groups are exogenously slightly different. Rather than through churning, their results arise through the dynamics of employer beliefs. In contrast with Cornell and Welch (1996), we assume no differences in the monitoring technology available for Black and White workers. In their paper, White employers can screen White workers more accurately than Black workers (or, in an extension, it is cheaper to monitor White workers) and are, therefore, more likely to hire a White worker.

Our empirical analysis begins with suggestive evidence that Black workers are more heavily supervised even in similar occupations, at least if they have no more than a high school education, a condition applying to the vast majority of Black workers during the period for which we have data on supervision.

Importantly for 'testability', our model has excess empirical content, predictions not known to be true or false before we developed the theory. First, it predicts that involuntary separations from employment will initially be higher for Black workers than for White workers, but that these hazards will converge with seniority. As seniority increases, it is more likely that workers have passed monitoring and are good matches with the firm. We test and largely confirm this previously untested prediction using the National Longitudinal Survey of Youth 1979 (Bureau of Labor Statistics, 2019). By two years of tenure, the magnitude of the gap has decreased substantially. This finding is robust to various sample selection decisions, approaches to smoothing the separation hazard functions, measures of seniority and proxies for involuntary separation, and strengthens with the inclusion of controls.

In addition, our model implies that high unobserved ability will have a larger effect on reducing unemployment and layoffs among Black than among White workers. Following a tradition dating from Farber and Gibbons (1996), we treat performance on the armed forces qualification test (AFQT), after controlling for observables, as unobserved by the market. We show that AFQT has a stronger negative effect on layoffs for Black than non-Black workers.

There are multiple equilibria in our model, a property it shares with models of rational stereotyping or self-confirming expectations (Coate and Loury, 1993). However, in our model, a group that begins with a low level of skills for which only the bad (monitoring) equilibrium exists will remain in that equilibrium even if its skill level rises to a level consistent with the existence of both the good and bad equilibria. Even if Black workers are, on average, more skilled than White workers, Whites can be in the good steady state and Black workers in the bad steady state because of a history of lower access to schooling and other human capital investments. Equalising the human capital that Black and White workers bring to the labour market may be insufficient to equalise labour market outcomes. In contrast, in self-confirming expectation models, if we could convince Black workers to invest in themselves and employers that they are investing, we would transition to the good equilibrium.

There is abundant evidence that Black workers face lower wages and longer unemployment duration than White workers. Moreover, these disparities are less prevalent (and perhaps, in some cases, non-existent) for the most skilled workers as measured by education or performance on the AFQT. While there are a plethora of models intended to explain wage or unemployment differentials, none addresses both and their relation to skill.[5] Since in our model newly hired Black workers are on average less productive than White workers, their wages are lower and firms that expect to hire Black workers anticipate less profit from a vacancy and therefore offer fewer jobs. Consequently, Black workers have longer unemployment durations.

We believe that the broad implications of our model can be derived through a variety of formalisations. The key elements common to these are:

(*i*) that a worker's productivity at different firms is correlated,
(*ii*) that workers cannot or do not signal their ability and that they can, at least imperfectly, hide their employment histories,[6]
(*iii*) that firms must, therefore, to some degree, statistically infer worker ability,
(*iv*) that further information about match productivity arrives during production and is either costly, imperfect or both, and

[5] Many models (e.g., Becker, 1971; Aigner and Cain, 1977; Lundberg and Startz, 1983; Lang, 1986; Moro and Norman, 2004; Bjerk, 2008; Charles and Guryan, 2008; Lang and Manove, 2011) assume market clearing and, therefore, cannot address unemployment patterns. Search models (e.g., Black, 1995; Rosen, 1997; Bowlus and Eckstein, 2002; Lang and Manove, 2003; Lang *et al.*, 2005) can explain unemployment differentials, but assume otherwise homogeneous workers and thus cannot address wage differentials at different skill levels. Peski and Szentes (2013) treated wages as exogenous. Coate and Loury (1993), Cornell and Welch (1996), Fryer (2007) and Bardhi *et al.* (2019) treated wages in specific jobs as exogenous, but endogenised the assignment of workers of different races to jobs so that, in this sense, they did address wages. In general, discrimination models have not addressed employment or job duration. See the review in Lang and Lehmann (2012).

[6] In particular, they must sometimes be able to omit or mischaracterise prior bad matches and misreport their time in the market.

(*v*) that this information, if obtained, may affect retention, so that firm behaviour affects the average unemployed worker's ability.

Our desire for a theoretically rigorous model of wage setting in a dynamic framework with asymmetric information drives the details of our formal model. Firms and workers bargain over wages and use a costly monitoring technology to assess the quality of the match, which is correlated with the worker's underlying type. Alternative models yielding the same intuitions include one in which signals are free and bad realisations cause Black workers, but not White workers, to be fired.

Therefore, use of the monitoring technology depends on the firm's prior: if the belief that a worker is well matched is sufficiently high or sufficiently low, it will not be worth investing resources to determine match quality. However, if the cost of determining the match quality is not too high, there will be an intermediate range at which this investment is worthwhile. Firm beliefs about Black, but not White, workers fall in this region. Consequently, they are subject to heightened scrutiny and are more likely to be found to be a poor match and fired. The increased scrutiny ensures that the pool of unemployed Black workers has a higher proportion of workers revealed as a poor match at one or more previous jobs. And, therefore, employers' expectation that Black workers are more likely to be poor matches is correct in equilibrium. This, nested in a search model, generates the empirical predictions discussed above.[7]

This churning equilibrium is hard to escape. Since education is observable, increased educational attainment might eliminate discrimination for those Black workers who acquire sufficient education. However, those with low education will still suffer discrimination relative to low-education White workers. Policy operating on unobservable skills—such as upgrading schools—is also unlikely to resolve the problem. Only if the skill level of Black workers is raised sufficiently above that of White workers[8] does the bad equilibrium cease to exist and White and Black workers receive similar treatment.

1. The Model

1.1. *Setup*

There are two worker groups, 'Black' and 'White'. Race is observable by the worker and employers, but does not have any direct impact on production.

At all times, a steady flow of new workers is born into each population group.[9] A proportion $g \in (0, 1)$ of new workers are type α, for whom every job is a good match.[10] The rest, referred to as type β, have probability $\beta \in (0, 1)$ of being a good match at any particular job. The probability of a worker being good at a job, conditional on her type, is independent across jobs. Worker type is private to the worker. Workers begin their lives unemployed. Without conditioning on type, the ex ante probability that a new worker is good at a particular job is

$$\theta_0 = g + (1 - g)\beta.$$

[7] Note that our model abstracts from moral hazard and that performance is observed objectively. MacLeod (2003) developed an interesting model in which biassed subjective assessments interact with moral hazard concerns.

[8] Technically, if the proportion of good workers is sufficiently high.

[9] We do not allow for death, but could do so at the cost of a little added complexity.

[10] Having type-α workers perform well at every job is not essential to the argument, but simplifies the presentation significantly.

Employers cannot directly observe worker type or employment history,[11] but can instead draw statistical inferences from race.

1.2. *Match Quality*

Production, the payment of wages and the use of the monitoring technology occur in continuous time using a common discount rate r.

Workers can be either well suited to a task (a 'good' match), producing q per unit time, or ill-suited (a 'bad' match), producing expected output $q - \lambda c$ per unit time. We can interpret the lower productivity of bad workers as errors or missed opportunities, each costing the firm c, that arrive at a constant rate λ. Under this interpretation, opportunities for error are also opportunities to learn the quality of the match, as well-matched workers are observed to avoid errors.[12]

The employer does not know the match quality without monitoring. During production, the firm may use a technology that may produce a fully informative signal about match quality. If the signal shows the match to be bad, the firm may terminate it immediately, receiving 0. The firm also has the option to cease monitoring while keeping the worker, for instance after a signal reveals the match to be good. Each moment, the firm chooses whether or not to monitor.

In keeping with the opportunity-for-error interpretation, we assume that the signal arrives at a constant rate λ. The monitoring technology costs b per unit time, so that the expected cost of information is $\int_0^\infty be^{-\lambda t}dt = b/\lambda$ and its expected discounted cost is $\int_0^\infty (e^{-rt}b)e^{-\lambda t}dt = b/(\lambda + r)$. As the signal arrives at the same rate regardless of match quality, this cost is unaffected by the firm's beliefs. The principal benefit of a signal whose arrival is exponentially distributed, rather than deterministic, is that it makes the employment survival function more realistic. In addition, it allows for a certain stationarity in the model: so long as no signal has arrived, the underlying incentives do not change. Optimally, following a signal that reveals the match to be good, the firm ceases monitoring, and the match continues indefinitely.

Conversely, for monitoring ever to be useful, matches revealed to be bad must separate. A sufficient condition for this is that $q - \lambda c < 0$. Additionally, we intend that β workers will not be willing to reveal their type in bargaining. To this end, we make the sufficient and simple assumption that such a match is unproductive, regardless of the monitoring choice:

$$\max \left\{ q - (1-\beta)\lambda c,\ \beta\frac{q}{r} + (1-\beta)\frac{q-\lambda c}{\lambda + r} - \frac{b}{\lambda + r} \right\} \leq 0. \tag{C1}$$

It is *much* stronger than necessary. In general, it is sufficient that any wage at which a firm would knowingly hire a β worker is low enough that the worker would rather reject it in order to rematch at a higher (pooling) wage. Assumption (C1) ensures that such separation in search of a new match is beneficial regardless of the expected duration of unemployment.

1.3. *Job Search*

When a worker is born or her match is terminated, she becomes unemployed. Unemployed workers are stochastically matched to firms, which occurs at a constant hazard μ. For the moment,

[11] At a more informal level, we believe that workers have some ability to hide their employment history and that they will not report information speaking to their own low ability.

[12] Alternatively, we could assume that the flows are $q - d$ and q with $d \equiv \lambda c$ and that λ is the arrival rate of opportunities to measure the flows.

we treat this rate as exogenous; it will be endogenised in Section 3.4 to address unemployment duration. When a match dissolves, transfers cease and the worker becomes unemployed. A firm does not recoup a vacancy and therefore receives a payoff of 0 on termination.[13]

In the unemployed state, workers merely search for new jobs; we normalise the flow utility from this state to 0. The value from unemployment is thus simply the appropriately discounted expected utility from job finding and is invariant to history. The expected discount on job finding is $\int_0^\infty e^{-rt}\mu e^{-\mu t}dt = \mu/(\mu + r)$; the value of a new job will depend on the equilibrium. We denote the expected present discounted value of future wages for an unemployed worker as U_θ^α for type-α workers and U_θ^β for type-β workers, when employers believe new matches to be good with probability θ. These will be constant in steady state.

1.4. *Wage Setting*

Given the asymmetry in information between the worker, who knows her type, and the firm, the Nash bargaining model is unusable and the Rubinstein (1982) one suffers from a multiplicity of equilibria. If a β worker does not want to reveal her type, as follows from our assumptions, then the β worker will have to bargain as if she were an α worker. Since in this case the firm cannot distinguish with which type it is bargaining, it should act as if it were bargaining with a random draw from the unemployment pool. Thus, an intuitively appealing solution is the outcome that would be reached in Nash bargaining between a firm calculating its rents on the assumption of a random draw from the pool and a worker calculating her rents as if she were an α worker.

We posit a simple wage bargaining model akin to Lauermann and Wolinsky (2016) that produces this outcome, albeit only in expectation. When a worker and firm meet, a wage offer w is randomly drawn from some distribution F. They then simultaneously choose whether to accept or reject the offer. If either rejects the offer, the match is dissolved; the firm receives 0 while the worker searches for the next firm. If both parties accept the offer, production proceeds at that wage. We are looking for perfect Bayesian equilibria in which neither party uses a weakly dominated strategy.[14] Using a randomly drawn take-it-or-leave-it offer allows us to escape both multiplicity of equilibria enforced by unreasonable off-path beliefs if the worker can make offers, and the Diamond paradox if only the firm makes offers.[15]

To ensure that the wage process does not end in disagreement in equilibrium, we assume that only agreeable wages are proposed. Specifically, we assume that F is a uniform distribution on the set of wages the firm and worker would both accept; thus, wage negotiation is as though an arbitrator proposes any wage on the contract curve with equal likelihood. Crucially, this assumption guarantees that, once we endogenise the job-finding rate, only lower demand for Black workers, and not the bargaining process, causes disparities in the formation of new matches.[16] Fortuitously, the assumption also results in simple solutions.

[13] This occurs naturally due to free entry when vacancy creation is endogenised; see Section 3.4.

[14] This is needed to rule out equilibria where both parties reject mutually acceptable wages.

[15] Earlier versions of this paper used an alternating-offer bargaining model with off-path belief restrictions and derived an equivalent set of theoretical results. The somewhat artificial nature of the current wage-setting structure dramatically simplifies the presentation without fundamentally changing the results.

[16] This makes F an equilibrium object, as the acceptability of wages depends on F, but the solution is unique in steady state. We could instead assume that F is uniform on $[0, q]$, but then unacceptable wages would be encountered, and the probability of disagreeable wages would vary between matches with Black and White workers. Note that F does not depend on worker type as, due to assumption (C1), there are no wages the firm and β workers would accept, but α workers would not. As equilibrium acceptable wages will form an interval, we could, instead of a uniform, use any distribution with connected, compact support by scaling it to the acceptable wage interval.

Jointly, our assumptions will ensure that every match will find a mutually acceptable wage, that equilibrium in steady state will be unique, that wages are uniformly distributed over the contract curve and that they are on average equal to the equal-weight Nash bargaining solution (between a firm with beliefs given by θ and an α worker), despite the asymmetric information.

1.5. *Steady State*

A steady state of a labour market is a mass of α job seekers, a mass of β job seekers and a mass of monitored β workers along with equilibrium firm and worker wage acceptance and monitoring strategies that make these populations constant over time. There is one steady state in which all employees are monitored until match quality is revealed, and one in which no monitoring occurs at all.[17]

Consider the case where no employees are monitored: the White labour market. Matches never deteriorate, and, therefore, the only source of job seekers is newly born workers. In this scenario, a firm just matched with a worker believes that the worker's probability of being of type α is the population prevalence g; the chance of a White job seeker being good at a job to which he is matched is therefore

$$\theta_W = \theta_0 = g + (1 - g)\beta.$$

Now suppose that all newly hired Black employees are monitored, and all bad matches are terminated. Newly matched Black workers will be worse than average. Surprisingly, the steady-state new match quality θ_B of this process does not depend on the rate of information λ, the worker matching rate μ or the rate at which new workers enter the market. This is an artefact of the assumption that workers are infinitely lived.[18] Simply, the quality of every match with a Black worker is eventually revealed, and bad matches are terminated. Therefore, every Black worker will rematch until they enter a good match. Thus, it takes one match for a Black α worker to exit unemployment forever, but an average of $1/\beta$ matches for a Black β worker. Consequently, Black βs will be overrepresented in the unemployed pool by a factor of $1/\beta$. Thus, the chance a new match is good is

$$\theta_B = \frac{g}{g + (1 - g)/\beta} \times 1 + \frac{(1 - g)/\beta}{g + (1 - g)/\beta} \times \beta = \frac{1}{g + (1 - g)/\beta}.$$

Our first lemma formalises this result.

LEMMA 1. *The probability a newly hired Black worker is in a good match is*

$$\theta_B = \frac{1}{g + (1 - g)/\beta} < \theta_W.$$

PROOF. See Online Appendix A.1. □

Therefore, although monitoring may be individually prudent for each firm, it creates a negative externality by feeding a stream of workers who are worse than the population average (i.e., containing more β types) back into the job-seeker pool.

[17] Technically, there is a third steady state in which firms sometimes monitor and sometimes do not, but it is fragile for reasonable values of the matching speed μ.

[18] In a model in which workers do not live forever, the steady-state expressions would be decidedly less elegant. On the other hand, in such a model we could allow for (slower) learning even in the absence of monitoring at the cost of some complexity. We interpret this as robustness to some kinds of endogenous monitoring intensity.

1.6. *Parametric Assumptions*

Now we impose certain restrictions on the joint values of parameters sufficient to ensure the existence of both steady states.

For an equilibrium with no monitoring to exist for White workers, we want to assume that monitoring costs are not too low. For monitoring not to be optimal, the instantaneous monitoring cost must not be worth paying to detect bad matches, accounting for the fact that the cost must be recouped on the surviving fraction of workers and thus

$$\underbrace{\frac{b}{\lambda}}_{\text{Monitoring cost}} > \underbrace{(1-\theta_W)\frac{\lambda c}{r}}_{\text{Reduction in losses to errors}} \cdot \underbrace{\theta_W}_{\text{Proportion of remaining workers}} . \tag{C2}$$

For the Black labour market, antisymmetrically to (C2), we posit that 'monitoring costs must not be too high'. We want to ensure that all new Black employees will be monitored in equilibrium. If there were no variation in Black workers' wages, the relevant condition would simply be

$$\underbrace{\frac{b}{\lambda}}_{\text{Monitoring cost}} < \underbrace{(1-\theta_B)\frac{\lambda c}{r}}_{\text{Reduction in losses to errors}} \cdot \underbrace{\theta_B}_{\text{Proportion of remaining workers}} . \tag{1}$$

However, as the monitoring decision is increasing in the wage, for all Black workers to be monitored, a condition is needed at the lowest wage in that market. As the lowest wage in the market depends on the speed at which workers match rather than simply the firm's break-even point, the relevant expression is a bit more complex:

$$\frac{b}{\lambda} < \underbrace{(1-\theta_B)\frac{\lambda c}{r}\theta_B}_{\text{Term from (1)}} - \underbrace{2(1-\theta_B)\frac{q-(1-\theta_B)\lambda c}{\mu+2r}}_{\text{Term due to search friction and wage variation}} . \tag{C3}$$

As the matching frictions vanish ($\mu \to \infty$), (C3) becomes (1).[19]

Thus, our condition stipulates that θ_B, the belief about the average ability in the Black unemployed pool, is sufficiently low that the firm monitors at all equilibrium wages. Strictness of the inequality ensures stability of the resulting steady state.

Finally, for labour markets to exist at all, it must be that workers can, in expectation, be gainfully employed. A sufficient condition for workers from both labour markets to be employable is that the expected product of workers drawn from the Black unemployed pool who are never monitored is positive:

$$q-(1-\theta_B)\lambda c > 0. \tag{C4}$$

2. Solution

First, we use the model's properties to characterise the firm's and worker's strategies. The main intuition behind the following result is that the firm is more willing to monitor if the bad matches terminated by monitoring are costlier, due to higher wages.

[19] To see how this inequality is used in proving proposition (C3), see Online Appendix A.4.

LEMMA 2. *The firm's monitoring decision is increasing in the wage w and decreasing in its belief about match quality θ.*

PROOF. See Online Appendix A.2. □

Our next result shows that the wages acceptable to both the firm and the workers form an interval.

LEMMA 3. *For a labour market with expected match quality θ, there is an interval of wages, $[\underline{w}_\theta, \overline{w}_\theta]$, the worker and firm both accept.*[20]

PROOF. See Online Appendix A.3. □

An intervalic structure for the wages in each market will simplify analysis significantly. From Lemma 3, the mutually acceptable wages are an interval $[rU^\alpha_\theta, \overline{w}_\theta] = [\underline{w}_\theta, \overline{w}_\theta]$. As α workers never separate once they find a job, the lowest wage is equal to the expected wage they would get at another firm, adjusted for search time, i.e., $\underline{w}_\theta = [\mu/(\mu + r)] \int_0^q w\, dF = \mu[0.5\underline{w}_\theta + 0.5\overline{w}_\theta]/(\mu + r)$, so that

$$\underline{w}_\theta = \frac{\mu}{\mu + 2r}\overline{w}_\theta.$$

We now present the main theoretical results of the paper: existence and uniqueness of equilibria in the two markets that perpetuate their associated steady states.

2.1. *The Non-Monitored Market*

PROPOSITION 1. *Assuming that (C1)–(C4) hold, the White (non-churned) labour market has a unique solution where the monitoring technology is not used. The average wage in this market is*

$$w^{avg}_{\theta_W} = \frac{\mu + r}{\mu + 2r}[q - (1 - \theta_W)\lambda c].$$

PROOF. See Online Appendix A.4. □

The main intuition for the proposition comes from Lemma 2. Since the value of monitoring is increasing in w, for a non-monitoring solution, we need only check whether the firm chooses to monitor at the break-even wage $\overline{w}_{\theta_W}$. Also, (C2) ensures that monitoring does not occur at that wage. Since firms do not learn workers' types in this labour market, White workers' types have no effect on their lifetime wages.

2.2. *The Monitored Market*

Here, as workers are monitored, β workers sometimes face separation and therefore have a low outside option. However, they cannot accept low wages at which monitoring would not occur at beliefs θ_B without revealing their type; thus, such wages are not accepted by the firm. Therefore, this equilibrium is effectively a pooling one as well, despite the fact that β workers receive significantly lower utility than α workers.

[20] Incentives are weak at the interval's endpoints, but this is immaterial as F will put zero probability on them.

PROPOSITION 2. *Assuming that (C1)–(C4) hold, the Black (churned) labour market has a unique solution where the monitoring technology is used in every match. The average wage in this market is*

$$w_{\theta_B}^{avg} = \frac{\mu + r}{\mu + 2r}\left[q - \frac{r(\lambda c(1-\theta_B) + b)}{\lambda\theta_B + r}\right].$$

PROOF. See Online Appendix A.5. □

The intuition here again comes from Lemma 2, which tells us that the monitoring decision is increasing in w and therefore if monitoring occurs at $\underline{w}_\theta$ it occurs at all matches, and (C3), which ensures that this condition holds. As the equilibrium strategies induce monitoring at every equilibrium wage, employees who are revealed to be in bad matches separate from the firm. This sends only β workers back into the job-seeking pool, churning the market quality to θ_B. On the other hand, when monitoring produces a good signal, the firm ceases monitoring and the worker is employed indefinitely.[21]

3. Implications for Labour Markets

The previous sections establish conditions under which there are two distinct steady states of the labour market. This section compares labour market outcomes for workers in these steady states. We first discuss a prediction that has not previously been tested and then discuss the relation of our other predictions to known labour market regularities.

3.1. *Job Duration*

Absent monitoring, there is no new information to dissolve the match. Therefore, taken literally, the model implies no turnover in the White equilibrium. In contrast, with monitoring, some workers prove ill suited for the job and return to the unemployment pool. We interpret this as predicting that Black workers will have lower average job duration. Recall that workers who return to the unemployment pool are all type β. Therefore, turnover is even higher than if only new entrants were monitored. The model, again taken literally, implies that the separation hazard for Black workers is

$$h(t) = \frac{(1-\beta)(1-g)\lambda e^{-\lambda t}}{1-(1-\beta)(1-g)e^{-\lambda t}},$$

which is decreasing in t, the amount of time passed in the match.

Importantly, h declines with t and asymptotes to 0, the hazard rate for Whites. We expect this prediction to be robust to important real-world elements not addressed by the model. Whether the hazard rates actually converge is not something we are aware of the literature addressing and is the subject of our empirical investigation later in this paper.

[21] That the firm ceases monitoring only when the signal arrives is a consequence of the fact that beliefs are constant in the absence of a signal. This is because the signal's arrival rate does not depend on the match quality. If bad matches were revealed faster than good ones, beliefs would drift upwards in the absence of a signal, and the firm would eventually stop monitoring. For a small drift, our analysis would remain largely unchanged. If, however, good matches were revealed faster than bad ones, beliefs would drift downwards in the absence of a signal, even leading to termination. In this latter case, as a match with a higher wage would be terminated faster, workers would not always prefer a higher wage, making our wage determination model inapt.

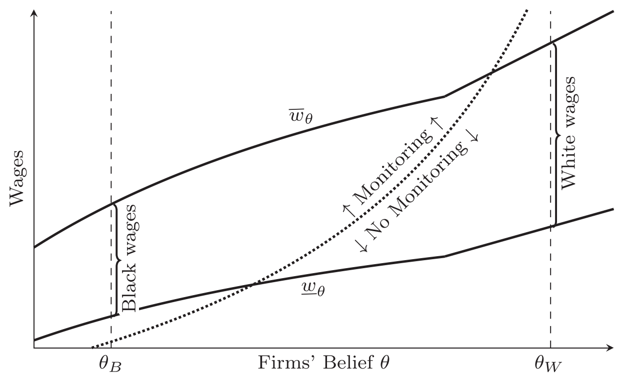

Fig. 2. *An Illustration in which the Black and White Wage Ranges Overlap.*

As our model abstracts from firm-to-firm hiring, we have no prediction regarding it. Although it may seem that firms would be out to poach Black workers with high seniority (that are likely to have passed monitoring),[22] adverse selection (with the worst workers more willing to leave) could unravel such effects. Still, our predictions are in terms of employer-initiated separations, not moves to better jobs. Therefore, in the empirical section, we treat spells that end in a job-to-job transition as censored and, in the main specification, treat all quits as censored.

We take this prediction one step further. Taken literally, our model implies that White workers are never laid off regardless of their type. In contrast, Black βs, but not Black αs, are sometimes laid off depending on whether they are good matches. We interpret this as a prediction that Black layoffs will be more responsive than White layoffs to a measure of unobserved worker quality.

3.2. *Wages*

As $w^{avg}_{\theta_B} < w^{avg}_{\theta_W}$, Black workers, on average, earn less than White ones. The highest wage firms are willing to pay is lower for Black than for White workers since the average quality of new hires is lower. The lowest wage Black workers are willing to accept is lower because they expect other employers to pay less, as well. Interestingly, because White workers are not monitored, their lifetime utility does not depend on type, and both types have higher utility than Black αs who, in turn, have higher lifetime utility than Black βs:

$$U^{\alpha}_{\theta_W} = U^{\beta}_{\theta_W} > U^{\alpha}_{\theta_B} > U^{\beta}_{\theta_B}.$$

Comparing utilities within type across the two groups, both types of Black workers are disadvantaged due to coming from a churned unemployed pool. This leads to lower wages and, when firms can direct searches as in Section 3.4, having to wait longer for matches. But Black β workers also suffer an additional consequence of being monitored more intensely—separation. This last group, therefore, suffers a combination of wage and retention discrimination.

For some parameter values, Black and White wage ranges overlap; for others, they do not. Figure 2 illustrates an example for which they do. Importantly, the average wage differential

[22] We show in Cavounidis and Lang (2015) that one can write a very similar model in which β workers always match badly, but monitoring can result in false positive good matches. Much of the analysis would remain unchanged in such a model. In such a model, Black workers would not eventually be better matched, on average, than White ones.

between Black and White workers is a consequence of differential monitoring, but not a requirement for it. If the (legal) requirement of equal wages for Black and White workers is effectively enforced, employers' incentives to monitor Black workers are stronger due to increased wages, and their incentives to monitor White workers are weaker due to reduced wages. Therefore, the range of parameter values consistent with discriminatory monitoring is broadened under such a requirement.

3.3. *Persistence of Discrimination*

A key feature of the churning mechanism in this paper is that deleterious steady states are persistent. We now discuss the difficulty of eliminating the bad steady state through policy aimed at upgrading the unobservable skills of Black workers (e.g., by improving schooling quality). The existence of a range of g values for which both steady states exist allows us to talk about *persistence* of the deleterious equilibrium.

Suppose now that rather than being identical, skill levels are $g_B \neq g_W$. Monitoring will persist as an equilibrium in the Black labour market until g_B rises above some critical level, while the no-monitoring equilibrium will exist in the White market provided that g_W remains above a lower critical level. In principle, we can have the Black workers in the bad equilibrium and the White workers in the good equilibrium provided that (C2) and (C3) hold for θ_W and θ_B calculated using g_W and g_B, respectively. Put simply, this means that discrimination in wages and monitoring (and therefore also separations) can continue even if Black workers are significantly better, on average, than White workers.

3.4. *Unemployment Duration*

We have so far treated unemployed workers' matching rate, μ, as exogenous. Making the standard assumption of free entry, we now allow firms to post and maintain vacancies at a cost k per unit time. When a firm creates a vacancy, it can direct its search. This can take several forms, most notably locating production operations in an area with specific population characteristics or advertising the vacancy in different areas and through different media. In general, a firm can target markets indexed by i where a proportion ρ_i of unemployed workers are White. The open vacancy cost k is invariant to this target choice. We assume that in each market i the bargaining equilibria and population group steady states break down along the discriminatory lines described so far.

Define φ as market tightness and let the worker job-finding rate function follow the commonly assumed form

$$\mu(\varphi) = m\varphi^{\gamma}$$

for constants $m > 0$ and $0 < \gamma < 1$. Note that if firms expect a match to be worth V, the free-entry level of φ in such a market sets

$$\frac{\mu(\varphi)}{\varphi}V - k = 0,$$

so that

$$\varphi = \left(\frac{Vm}{k}\right)^{1/(1-\gamma)}.$$

Therefore, φ is an increasing function of V.

Assuming that the parametric assumptions hold for the entire breadth of derived matching rates, we can now derive the free-entry equilibrium level of μ_{ρ_i} for each market i. When hiring from pool i, the firm's expected payoff from a new match is

$$V_i = \rho_i \frac{1}{\mu + 2r}[q - \lambda c(1 - \theta_W)] + (1 - \rho_i)\frac{1}{\mu + 2r}\left[q - \frac{r(\lambda c(1 - \theta_B) + b)}{\lambda \theta_B + r}\right].$$

The above expression is strictly increasing in ρ_i. Therefore, for the same μ, markets with more Black workers will have a lower expected payoff for a filled vacancy. Therefore, the free entry $\varphi(\rho_i)$ and $\mu(\varphi(\rho_i))$ are strictly increasing in ρ_i, so that workers searching for jobs in markets with a higher proportion of Black workers take longer, on average, to find employment. Online Appendix A.6 shows that this implies that an unemployed Black worker must wait longer for a match, on average.

3.5. *Extensions and Further Discussions*

3.5.1. *Wage renegotiation and evolution*

There are many ways to think of wage evolution in models with learning about match quality. For instance, Bose and Lang (2017) envisioned the firm and worker as splitting the instantaneous surplus according to a fixed rule. However, the presence of asymmetric information in our model complicates this. In line with Lazear (2009) and Postel-Vinay and Robin (2002), we choose to model wage renegotiation as responding to outside offers.

As in Postel-Vinay and Robin (2002), new wage offers come from firms with the same information as the incumbent; they know if monitoring concluded successfully (alternatively, the worker can credibly disclose a success). Consistent with our wage determination model, the outside offer is uniformly distributed between the worker's current wage and the wage at which the outside firm expects to break even, given the available information. The incumbent firm then matches this offer and keeps the worker.[23]

For White workers and Black workers whose monitoring had not yet concluded, this upper bound would be the same as the break-even wage at the incumbent firm. The reason for this is that the firm's belief about match quality is constant as long as no signal arrives. The outside firm would prefer to monitor a Black worker who is being monitored. If a Black worker had been revealed to be a good match in the past, on the other hand, this would be informative. However, being a good match at one job does not mean the worker would automatically be good at the next one! The probability of a Black worker being an α given a monitoring success is

$$\begin{aligned} P(\alpha \mid success) &= \frac{P(success \mid \alpha)P(\alpha)}{P(success)} \\ &= \frac{1 \cdot \{g/[g + (1 - g)/\beta]\}}{1 \cdot \{g/[g + (1 - g)/\beta]\} + \beta \cdot \{(1 - g)\beta/[g + (1 - g)/\beta]\}} \\ &= g, \end{aligned}$$

which is the same as that for a White worker who has never been monitored. Thus, an outside firm would offer a wage uniformly distributed between the old wage and the highest White wage.

[23] To model both transitions and within-firm wage evolution without adding firm heterogeneity, we could assume that the incumbent only sometimes has the opportunity to respond.

Allowing for renegotiation, our predictions regarding average wages by race would be preserved, conditional on any of age, experience or job spell duration.[24] However, Black workers would at all times see larger average increases in log wages—a steeper wage profile. This is because White workers and Black workers with a monitoring success both approach a wage of $q - \lambda c(1 - \theta_w)$ at a rate proportional to their distance from it, but Black workers begin at a lower starting wage. Bratsberg and Terrell (1998) found that the return to tenure is no lower and often higher for Black than for White male high school graduates, although this finding is somewhat sensitive to choice of estimation technique.

3.5.2. *Skill level and discrimination*

Furthermore, we can allow for observable heterogeneity among workers. If there are groups of workers for whom g is high, only the no-monitoring equilibrium will exist for these groups, regardless of race. This is also true at very low g and very low β (although we have assumed away this case to simplify the proofs). The first result is consistent with similar outcomes for Black and White workers with high levels of skill as measured by education (Lang and Manove, 2011). The latter is consistent with some evidence that the bottom of the labour market is similarly bad for Black and White workers. On the other hand, Lang and Manove (2011) found that the market learns the productivity of White, but not Black, high school dropouts. This is consistent with an equilibrium in which White unemployed dropouts are, on average, more skilled than Black unemployed dropouts and therefore in which White, but not Black, dropouts are monitored. Nevertheless, without additional, largely ad hoc assumptions, this story cannot account for the very high unemployment rate among Black dropouts.

3.5.3. *Changing screening and monitoring technology*

Autor and Scarborough (2008) examined the effect of bringing in a new screening process. They found that the screening process raised the employment duration of both Black and White workers with no noticeable effect on minority hiring. In our model, we can think of this technology as allowing the firm to screen for job match quality prior to employment, successfully detecting bad matches with some probability. This increases the proportion of hired Black workers who become permanent since some bad matches are not hired. If the screening mechanism is good enough, the firm will choose not to monitor the Black workers it hires, and all Black workers will be permanent. Formally, since all White workers are permanent in the absence of the screen, the screen does not affect this proportion. Informally, if poor matches are more likely to depart even without monitoring, then there will also be positive effects on White employment duration.[25] Similarly, Wozniak (2015) showed that drug testing increases Black employment and reduces the wage gap; we interpret this as confirming evidence for the notion that employers are more uncertain about the quality of Black workers, and therefore that Black workers benefit more from early resolution of such uncertainty.[26]

[24] Naturally, as workers are able to use evidence of a good match to negotiate higher wages, firms benefit less from monitoring. This means that they are less willing to monitor. We would thus need different parametric assumptions for both of our equilibria to exist, which we could phrase in terms of the arrival rate of opportunities for renegotiation to be slow enough.

[25] Formally, the model would have to be modified to ensure that some β workers are never perfectly matched and/or that some β workers are still in bad matches when they exit the labour force.

[26] Wozniak (2015) is not to be interpreted as evidence that monitoring is good for Black workers in the aggregate. As in the present paper, it can be beneficial on an individual level (as it allows good workers to get higher wages than otherwise); our model, however, shows that it can also create a worse externality.

We note that improved technology appears to have reduced monitoring costs. This is unambiguously good for Black workers who share the cost of being monitored. Unless the reduction shifts Whites into the monitoring equilibrium, they are unaffected by the cost reduction. However, if firms begin monitoring White workers, White αs and firms will initially be better off. Firms will be able to better screen their workers, and, as a consequence, can offer higher wages, which should make α workers better off as the monitoring does not put them at risk. On the other hand, β workers will generally be worse off. In a collective bargaining setting, the union might resist monitoring. The more interesting point is that, since monitoring creates an externality, it is easy to develop an example in which monitoring makes both types of workers and capital worse off in the long run.[27]

3.5.4. *Segregation*

It would be a clear violation of US civil rights law for firms to monitor Black, but not White workers. However, we show that Black workers, coming from a low-quality unemployment pool, will prefer to apply to jobs with monitoring while the opposite will be true for White workers, who come from an unchurned pool. Consequently, our results will go through, with necessary tighter parametric restrictions, provided workers and/or firms can target their searches. In this section, we make this point for a world with heterogeneous jobs, but it holds mutatis mutandis for homogeneous jobs.

So far, we have assumed that monitoring is available at all firms at an identical cost. But what if it is not? Would the equilibrium necessarily unwind as Black workers take jobs without monitoring, and the Black unemployed pool becomes less churned? In the simplest model, in which some firms simply cannot monitor, if matching is fast enough that Black workers can afford to avoid no-monitoring jobs, the answer is no.

Suppose that the probability of matching to a firm with no-monitoring technology is p and that the probability of matching to a firm with the usual technology is $(1 - p)$. Black workers receive lower average wages at no-monitoring jobs, so Black α workers (who lose nothing from being monitored) prefer to match with firms where monitoring is available. As Black βs take any wage the αs do, at any wage a non-monitoring firm and Black worker agree on in equilibrium, that firm's beliefs can at best be θ_B. Thus, the highest wage that could occur in such a match is $q - \lambda(1 - \theta_B)c$. For a Black α worker to take such a wage, it must be that they prefer it to their utility from searching until a monitoring firm is found, or

$$q - \lambda(1 - \theta_B)c \geq \frac{\mu(1 - p)}{\mu(1 - p) + 2r}\left[q - \frac{r(\lambda c(1 - \theta_B) + b)}{\lambda\theta_B + r}\right].$$

However, in the limit as μ goes to infinity, this is strictly ruled out by (C3). Therefore, if μ is high enough, Black α workers will not take no-monitoring jobs, so that if Black β workers took such jobs, they would be revealed (and the firm would thus reject). As a consequence, no Black workers take a no-monitoring job in equilibrium.

The takeaway is that Black αs *do not want* no-monitoring jobs, and will avoid them if they can wait for monitoring ones, thus forcing β workers to follow suit or be revealed. The fact that *monitoring is beneficial to individual Black α workers* is part of what makes the bad

[27] Suppose that g_0 is just sufficient to sustain a no-monitoring equilibrium. A small reduction in b puts the labour market into a monitoring equilibrium. Initially, α workers and firms would experience a slight gain, but the churning will wipe this out and more. Firms always make zero profit on vacancies, but if we allow for a distribution of vacancy costs then the rents earned by firms with low costs of creating vacancies will also fall.

equilibrium Black workers find themselves in so robust. So, rather than sorting to no-monitoring jobs and unchurning the unemployed pool, Black workers avoid such jobs, whereas White workers do not.

4. Empirics

4.1. *Are Black Workers Monitored More?*

Our theoretical analysis assumes that Black and White workers are in different equilibria and that, consequently, Black workers are more heavily monitored. Of course, it is possible that some Black workers end up in low-wage jobs where monitoring is unnecessary after employers observe their previous separation history. In addition, a more realistic model would have Black workers more heavily monitored than White workers, but would not predict that White workers are never monitored.

However, the spirit of the model is that Black workers are monitored more heavily, and therefore, we look for direct evidence related to this result. The evidence we have been able to find is minimal. We rely on a single question from the 1977 wave of the Panel Survey of Income Dynamics (PSID) that asks whether the respondent's supervisor checks his work 'several times a day, once a day, once a week, every few weeks, or less often than that' (Panel Study of Income Dynamics, 2020). We code the reported level of supervision from one to six, where six denotes supervision several times per day, and one denotes reporting not having a supervisor.

We do not know the nature of the supervision. Ideally, we would like a measure of supervision designed to assess the worker's quality rather than supervision to prevent shirking or other malfeasance. Neal (1993) used this variable to study differences in supervision focussed on the latter. Still, this is what appears to be available to us.

We have supervision data only on household heads who are actively employed or temporarily laid-off private-sector workers.[28] We exclude individuals who are supervisors.[29] We estimate the model by ordered probit and weight by the 1977 family weight. Early experimentation found only weak evidence of monitoring differences when we did not further restrict the sample and no evidence that Black workers with more than a high school education were monitored more frequently than their White counterparts. Perhaps, as suggested in Section 3.5.2, more educated workers are more likely to be good types and therefore educated workers of both types are in the no-monitoring equilibrium. Alternatively, the question may not be good at revealing monitoring differences among more skilled workers. Nevertheless, in the remainder of the paper, we focus on workers with no more than high school education.[30] Recall that Black workers with more than

[28] In a small number of cases, a non-household head answers about supervision of the household head at the head's job. We exclude these cases given concerns about measurement error. The PSID asks the employment survey questions to individuals with an employer and who 'are working now or are reasonably likely to return to work in the near future'. Thus, temporarily-laid-off workers should respond based on the job to which they soon expect to return.

[29] We exclude supervisors, as Black workers who supervise other workers are intuitively more likely to have passed monitoring. We also restrict our sample to respondents living in the United States who report a wage. Respondents are only asked to report a wage if they report being salaried or paid hourly. Respondents who replied 'other' or 'NA; Don't Know' to whether they were paid by the hour or salaried were not asked for their wage. The question about supervision is also asked separately to individuals who report working for someone else and being self-employed. We do not include these individuals as we cannot separately identify the occupation and industry for these two jobs.

[30] We note that earlier versions of this paper used workers regardless of education. In general, our results are strengthened by restricting the sample, but we did not think to use this restriction until we began to look at supervision directly.

Table 1. *Likelihood of Employer Monitoring by Race.*

$Y =$ level of supervision	(1)	(2)	(3)	(4)
Black	0.150	0.212	0.176	0.185
	(0.099)	(0.104)	(0.110)	(0.110)
Other race	−0.248	−0.225	−0.350	−0.334
	(0.176)	(0.199)	(0.208)	(0.211)
Completed education	≤12	≤12	≤12	≤12
Occupation, industry FEs	N	Y	Y	Y
Other controls	N	N	Y	Y
Ln(hourly wage)	N	N	N	Y
N	1,095	1,095	1,089	1,089

Notes: Robust SEs are given in parentheses. Estimates are from an ordered probit using data from the 1977 PSID. The dependent variable is the level of employer supervision. A value of six corresponds to the employer checking the individual's work several times per day, five to once a day, four to once a week, three to every few weeks, two to less often and one corresponds to no supervisor. Other controls include the highest grade completed, age, age squared, tenure, tenure squared and indicators for temporarily laid off, South, North Central, Northeast, salaried, male and union job. The sample includes household heads employed or temporarily laid off by private employers, who reported a wage and are not themselves supervisors. Observations are weighted by the family weights of the survey. See the text for details.

a high school education were a relatively elite group during this period. Only 11% of the Black workers in the sample (16% weighted) were in this group.

The first column of Table 1 presents the results with no controls; Black workers are more likely to report being monitored frequently. However, the coefficient is imprecisely estimated and significant at only the 0.1 level.[31] Including 12 occupation and 11 industry fixed effects (column (2); see Online Appendix B.1 for details) increases the coefficient, which is now significant at the 0.05 level. Note that adding these controls may be necessary or problematic. If monitoring costs varied among firms, we anticipate that Black workers would be more likely matched with jobs in which monitoring is relatively inexpensive (as in Section 3.5.4); on the other hand, if the sensitivity of output to skill varied between jobs, we anticipate that more Black workers would be matched to firms with low skill sensitivity. In the former case, the mechanisms in our model explain why Black workers are ending up in occupations and industries in which monitoring is more frequent. As a result, controlling for occupation and industry would obscure Black-White differences between occupations explained by our model. If Black workers end up in occupations and industries with low skill sensitivity because of the mechanisms of our model, they will be monitored less frequently. As a result, failing to control for occupation and industry will obscure Black-White differences within occupation and industry that may arise due to the mechanisms in our model. Adding these controls is a conservative approach.[32]

Column (3) includes additional controls for years of education, quadratics in tenure (truncated at the 99th percentile or 444 months) and age, male, union job, whether the worker is salaried, living in the Northeast, North Central or South of the United States (with West the omitted region) and whether the worker is temporarily laid off. Note that some of these controls are potentially endogenous in a fuller model. This yields a smaller coefficient that falls just short of significance at the 0.05 level. Finally, as a potentially better, but obviously endogenous control

[31] Throughout this subsection, we use one-tailed tests because we will not consider large negative *t*-statistics as evidence in favour of our alternative hypothesis, and we do not seek to identify evidence that rejects no difference against the alternative that Black workers are monitored less than White ones.

[32] For the coefficients on the controls, see Online Appendix Table B.1.

for worker skill, column (4) adds the natural log of the individual's wage, and the coefficient is again significant at the 0.05 level.[33]

Marginal effects in column (4) suggest that Black workers are 6.2 percentage points (pps) more likely to be monitored several times per day than White workers. Weighting by the 1977 family weight, 34.4% of White workers report being monitored several times per day. Thus, Black workers are 18% more likely than White workers to be monitored several times per day. Black workers are less likely to (*a*) report no supervisor (2.3 pps equivalent to 26% less, given 8.9% of White workers report no supervisor), (*b*) be monitored less often than every few weeks (2.6 pps equivalent to 15% less, given 16.9% of White workers report this level of supervision), (*c*) be monitored every few weeks (0.5 pps equivalent to 10% less, given 5% of White workers report this level of supervision) and (*d*) be monitored once a week (0.7 pps equivalent to 6.4% less, given 11% of White workers report this level of supervision). There is no difference in the likelihood of monitoring once a day.

Our model implies that monitoring should decline faster with tenure for Black workers than White workers. We tried including interactions between race and tenure and tenure squared. Unsurprisingly given the small sample, the results were uninformative.

4.2. *Unemployment, Race and AFQT*

Building on a literature starting with Farber and Gibbons (1996), we use AFQT to capture both unobservable and observable predictors of quality such as education. It is well known that Black individuals score lower than White individuals on the AFQT and increasingly well known that, conditional on AFQT, Black workers get more education than White workers. Thus, it would be surprising if Black unemployed workers did not have lower scores than their White counterparts on the AFQT. Nevertheless, for completeness, we verify this expectation as it confirms employers may find greater reason to monitor newly hired Black workers than White workers.

We use the National Longitudinal Survey of Youth 1979, a nationally representative sample of 12,686 individuals, 14–22 years old when first surveyed in 1979, with oversamples of Black, Hispanic and poor White individuals (Bureau of Labor Statistics, 2019). These individuals are surveyed annually through 1994, and biennially afterwards. We restrict ourselves to the non-Hispanic sample of Black and non-Black individuals and eliminate the poor and military oversamples, and use the most recent scaling of the AFQT. Respondent's current labour force status was recorded only in the waves through 1998 and again in 2006. We drop the waves after 1998 and, in each wave, individuals who were out of the labour force. We use the survey-week labour-force status from the NLSY, which asks about the respondent's main survey week activity.

AFQT is measured on an ordinal scale, at least with respect to worker productivity. Therefore, monotonic transformations of the scale can affect whether one group has higher AFQT unless the scores of one group are higher in the sense of first-order stochastic dominance (FOSD). Using the scale score rather than the percentile rank could change the ordering of groups. Consequently, we test for FOSD, or more precisely its absence, using the Kolmogorov–Smirnov test. Formally, this tests whether we can reject the null hypotheses that the distribution of group W(hite) dominates that of group B(lack) and that B dominates W. If we cannot reject either hypothesis, we conclude that we cannot reject that the distributions are equal. Importantly, rejecting that B dominates W,

[33] Our model implies that White workers earn more than Black workers. Thus, conditional on wage, Black workers should be more skilled than White workers on average. If more skilled workers are less likely to be supervised, controlling for wage should underestimate the true difference.

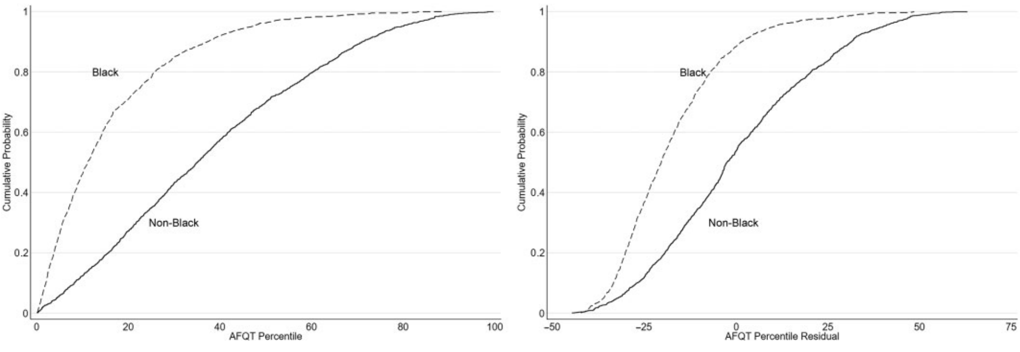

Fig. 3. *Distribution of Ever-Unemployed Workers' AFQT by Race.*
Notes: Left: CDF of AFQT percentile. Right: CDF of residual from regressing AFQT percentile on education, SMSA residency, age and region. Sample excludes Hispanic individuals.

but not that W dominates B does not allow us to *accept* the null that W dominates B. However, when combined with visual evidence suggesting stochastic dominance, we conclude that W dominates B.

The Kolmogorov–Smirnov test requires dividing the sample into two groups, such as Black and White workers. When comparing employed and unemployed workers, some sample members may be employed in some years and unemployed in others. Our primary analysis divides workers between those unemployed at the time of any interview and those who never report being unemployed, whom we informally call 'always employed' although they may have been out of the labour force at some interview.[34] We define education as the highest educational attainment the respondent reports in any year. We only retain respondents who never report education beyond high school.

Ever unemployed non-Black individuals have higher AFQT scores than ever unemployed Black individuals both visually—seen in Figure 3—and statistically. We can reject that Black scores dominate non-Black scores at any conventional level while the test statistic for the opposite hypothesis is < 0.0005. We also residualised AFQT using a regression of AFQT on education, a dummy for whether the individual lived in an SMSA when the AFQT was administered, their age in 1979 and three region dummies. Among those ever experiencing unemployment on an interview date, non-Black individuals have higher residualised AFQT scores than Black individuals, both visually and statistically.

Our model also implies that unemployed Black workers should be more adversely selected than their White counterparts. Unfortunately, we are unable to determine whether this prediction is empirically valid. For the reasons discussed in Bond and Lang (2013), it turns out that which group is more adversely selected depends on our choice of mapping from AFQT score to ability.[35]

[34] In principle, we could use multiple observations per individual and split the sample by whether the individual was employed or unemployed in a given year. However, the Kolmogorov–Smirnov test would then not have a standard distribution. While the figures show the distributions with weights, we are forced to perform the formal test on unweighted data. Fortunately, the correlation within race between AFQT and the sampling weight is small. The plotted CDFs are virtually indistinguishable.

[35] Relative to Bond and Lang, we have an additional degree of freedom in choosing our metric for 'more adversely selected'. We can choose the percentile below which we define someone to be a β. We can also choose whether our metric for 'more adversely selected' is the difference between Blacks and Whites in the difference between the unemployed and

4.3. *Initial Black versus Non-Black Gaps in the Layoff Hazard, and Convergence over Time*

We now test the model's prediction that the layoff hazard is initially higher for Black workers, but converges to that for White workers. To our knowledge, this prediction has not previously been tested. The model also predicts longer unemployment durations and lower lifetime incomes for Black relative to White individuals. These are known to be strongly empirically supported (see Lang and Lehmann, 2012). As discussed above, there is some evidence that the return to tenure is higher for Black workers than for White workers, as an extension of our model predicts.

Ideally, we would restrict separations to terminations for cause. While both layoffs and firings are involuntary separations, layoffs are technically not for cause. However, individuals may not respond with these distinctions in mind or may not want to respond that they were fired. Additionally, employers may choose to include some individuals in a layoff, rather than terminate them for cause.[36] For these reasons, we consider the hazard of being fired or laid off.

Focussing only on firings produces somewhat better results. Nevertheless, we show this as the robustness check we intended it to be. We show in the Online Appendix that layoffs due to plant closings do not exhibit the pattern we predict for firings.[37] Our results may underestimate the racial gap by including some exogenous layoffs. Furthermore, many separations reported as quits may be induced quits. If this is more common for Black workers, this will also underestimate the racial gap in firing/layoffs. As a robustness check, we show in the Online Appendix that treating quits into non-employment as layoffs does not substantially change our results.

4.3.1. *Methods*

We estimate the layoff hazard using the first full-time spell of each individual at each employer. We censor spells ending for any reason other than the employee being fired or laid off. For these censored spells, we assert that we do not know when the spell would have ended in a layoff. We use both non-parametric and semiparametric survival analysis methods to estimate the hazard.

First, using standard techniques, we calculate the hazard over time intervals with the intervals large enough not to require further smoothing. For each non-overlapping time interval, $(t_{j-1}, t_j],\ j = 1, \ldots, k+1$, we obtain the number of employment spells at the start, the number of spells ending in a layoff (failures) over the interval and the number of spells ending, but not in a layoff (censored). A conventional way of calculating the hazard in this setting is to assume that censoring and death times are uniformly distributed within each interval. The hazard at the midpoint m for each non-overlapping interval is then

$$\hat{h}(t_{mj}) = \frac{d_j}{(t_j - t_{j-1})(Y_j - d_j/2)}.$$

employed in probability of being a β, the difference in the odds ratio or the difference in the log odds ratio. These choices lead to very different conclusions. Choosing our scale and metric simply leaves us with too many degrees of freedom.

[36] Oyer and Schaefer (2000) found evidence consistent with firms substituting towards layoffs and away from firings for Black men relative to White men following the Civil Rights Act of 1991. The Civil Rights Act of 1991 increased the expected litigation costs of discharging employees protected by the legislation.

[37] We show the hazard for spells that we know end in layoffs due to plant closings, reported as a separate category starting in 1984. We do not treat these as layoffs in our main analysis. This exercise is underpowered, with approximately 850 spells ending for this reason. We treat individuals reporting before 1984 that their spells end due to the combined category of layoffs, plant closing, or end of temporary or seasonal jobs as censored in this robustness analysis (Online Appendix Figure B2(k) and (l)). Restricting to spells starting in 1983 or after yields results for this robustness exercise with a similar pattern.

The variable Y_j is the number of spells at the start of the interval minus half of the spells censored over the interval, and d_j is the number of failures over the interval (Klein and Moeschberger, 2003).[38]

We use intervals of 26 weeks through durations of 520 weeks. After this point, 26-week intervals no longer include at least one Black worker who was laid off and at least one non-Black worker who was laid off, and so we use intervals of 39 weeks. This facilitates comparison of Black and non-Black worker hazards in our subsequent Cox analysis, as these models provide estimates only at failure times.[39] We calculate the hazard separately over these intervals for Black and non-Black workers. We obtain confidence intervals based on the estimated SD of the hazard function at the midpoint of interval j, using the property that the number of failures in the interval is a binomial random variable.[40]

This non-parametric method does not allow controlling for covariates. So we additionally plot baseline hazard functions for Black and non-Black workers from a Cox proportional hazard model stratified by race. The stratified Cox model allows for different baseline hazard functions for Black and non-Black workers, rather than assuming that their baseline hazards are proportional. As in the traditional Cox model, we constrain the coefficients on the covariates to be the same for Black and non-Black workers. We use the time intervals defined above as our measures of time so that the baseline estimates do not require further smoothing.[41] The baseline contributions we obtain from this model are the same as the Nelson–Aalen contributions in the case of no covariates, using the week intervals as a measure of time.[42]

Specifically, we estimate

$$h(t \mid W, \mathbf{Z}) = h_W(t) \exp(\mathbf{Z}\boldsymbol{\gamma}).$$

The variable W is an indicator for whether the individual is non-Black and $\mathbf{Z}$ includes highest grade completed, indicators for geographic region (Northeast, North Central, South; West is omitted), whether the individual lived in an urban area, age, occupation and industry fixed effects, and year fixed effects, all measured at the start of the spell, and the AFQT percentile. As described in Online Appendix B.2.4, we classify occupation into 18 groups and industry into 14 groups.

Baseline hazard contributions are functions of the estimated coefficients from a Cox model. As a result, we obtain the variance of the difference in the Black and non-Black hazards at each failure time based on a non-parametric bootstrap, and 10,000 bootstrap samples. We then use these variances to construct confidence intervals for the difference between the hazard for Black and non-Black workers. As we describe in detail in Online Appendix B.2.5, for each bootstrap sample, we estimate the Cox model and then follow Kalbfleisch and Prentice (2002) to obtain

[38] This is referred to in the literature as the life-table method. We note that without the adjustments for the timing of failures and censorings, this is simply the fraction of spells still ongoing at the start of the interval that fail during the interval.

[39] These 39-week intervals include at least one Black and one non-Black worker failure through durations of 793 weeks (15.25 years). While we use all durations for estimation, our figures show hazard estimates through 793 weeks.

[40] The formula for the estimated SD of the hazard comes from Klein and Moeschberger (2003) and is used in STATA. Gehan (1969) derived a similar formula.

[41] Using time intervals rather than week as a measure of time creates more instances of failures occurring at the same 'time' since time is now a larger unit. There are several methods for dealing with these ties, all requiring assumptions about the timing of these failures. We present results using the Breslow approximation, one of the conventional methods, and the STATA default. This is based on the assumption that the subjects failed at different times, but we do not know the order.

[42] There are several estimators of the baseline hazard rate in a proportional hazard model. We use the estimator from Kalbfleisch and Prentice (2002), also the default in STATA.

the baseline hazard contributions at each failure time by solving for the maximum likelihood estimates using an iterative procedure.

In a robustness check, we determine the hazards at each week, rather than for an interval of weeks, and then smooth using a kernel smoother and local linear smoothing. These methods require choices of kernels and bandwidths and, in the former case, an approach to addressing bias in the boundary regions.

4.3.2. *Data*

We use the NLSY79 to test the model's prediction that the layoff hazard is initially higher for Black workers, but eventually converges to that of non-Black workers. We construct job spells using the Employer History Roster, which greatly facilitates linking job spells across survey years, by assigning each job a unique identification number consistent across surveys.[43] We define job spells as the first full-time spell with each employer, defining full time as at least 30 hours per week. For each survey year in which an individual reported employment at a given employer, we collect the start and end weeks of employment with that employer reported in the survey.

We construct the total length of the job spell by grouping all consecutive full-time spells at the employer across survey years. We treat gaps at the same employer of less than or equal to 26 weeks as continuations of the same job spell at the employer, but subtract the length of the gap from the duration.[44] We evaluate whether the individual is fired or laid off only at the end of the linked spell. For robustness, we do not link non-continuous reported spells at the same employer. We focus on the layoff hazard, which we define as the hazard of a job ending due to the employee being fired or laid off.[45]

Our sample includes non-Hispanic individuals who had obtained no more than a high school degree at the start of the job spell, consistent with our earlier analysis.[46] We exclude spells in which the worker ever reports self-employment or working for a family business. As we describe in Online Appendix B.2.1, we further exclude individuals with missing start or end weeks for any full-time spell, and individuals with full-time spells that end before they begin.

Because the survey is conducted every two years starting in 1994, we do not know the values of some of the control variables in some years and must impute their values from adjacent years. Online Appendix B.2.2 describes these imputations in detail. To avoid excluding individuals with missing covariates, we include an indicator for whether the individual is missing the covariate, and set the covariate to zero. As described in Online Appendix B.2.4, we convert all occupation

[43] Information for jobs six through ten reported in some of the early survey years may not have been added by NLSY to the roster due to difficulty recovering these data. This is unlikely to have a large impact on the results given that these jobs are a small proportion of those ever reported, for a small proportion of individuals (National Longitudinal Survey of Youth, 2019).

[44] The individual could report multiple job spells at the same employer, reporting a start and end week for each span. Additionally, for a given spell, the individual can report a within-job gap at the employer and the start and end weeks of that gap. We exclude spells in which the individual reported a gap of more than 26 weeks within the given start and end weeks they reported at the employer. The reported reasons for these within-job gaps make it difficult to identify whether the gap was due to the individual being fired or laid off, and so we exclude these spells.

[45] From 1979 through 1983, the NLSY groups together 'layoff, plant closed, or end of temporary or seasonal job' as one reason for the job ending. Starting in 1984, these three categories are separated, and we treat layoffs as failures and the other two categories within the original group as censored. NLSY groups together fired and discharged as a reason for leaving the job. For simplicity, we refer to this category as a firing.

[46] We show results for individuals with more than a high school degree at the spell's start for completeness. As above, we exclude the poor and military oversamples and include all surveys through 2010. We use the NLSY racial/ethnic cohort coding from the 1978 screener interview, which codes individuals as Hispanic; Black; or non-Black, non-Hispanic.

Table 2. *Summary Statistics.*

	Non-Black	Black
Spells ends in layoff/firing	0.21	0.23
	[0.41]	[0.42]
Spell duration	91.5	82.4
	[178.8]	[158.7]
Male	0.55	0.61
	[0.5]	[0.49]
Age at spell start	25.8	27.3
	[8.01]	[7.9]
Highest grade completed at spell start	11.3	11.4
	[1.23]	[1.11]
AFQT (percentile)	42.3	18.6
	[25.2]	[17.5]
Urban location at spell start	0.72	0.83
	[0.45]	[0.37]
Spells per person	5.65	6.26
	[5.17]	[5.28]
Occupation: managerial and professional	0.06	0.04
	[0.24]	[0.19]
Occupation: technical, sales, administrative	0.22	0.18
	[0.41]	[0.38]
Occupation: service	0.18	0.25
	[0.38]	[0.43]
Occupation: precision production, craft and repairers	0.13	0.1
	[0.34]	[0.3]
Occupation: operatives and labourers	0.23	0.29
	[0.42]	[0.45]
Spells at risk of ending in non-employment at		
200 weeks	2,298	1,406
400 weeks	1,071	578
600 weeks	587	288
800 weeks	331	160
1,000 weeks	197	99
Total spells	20,140	13,674

Notes: SDs are given in brackets. Sample excludes Hispanic workers.

and industry codes, which vary across survey year, to 1990 census occupation or industry codes, using IPUMS data (Ruggles *et al.*, 2022).

Table 2 shows summary statistics for nearly 34,000 job spells, which are the first full-time job spells at each employer for individuals in the sample. There are 20,140 job spells for non-Black workers and nearly 13,700 for Black workers. For non-Black workers, the average spell duration is 91.5 weeks, while for Black workers, the average spell duration is 82.4 weeks, though these are underestimated due to censoring. The table shows other differences between the average non-Black and Black job spells, including the worker's age, education, AFQT, urban location, occupation and region. Importantly, the stratified Cox proportional hazard models will include these as covariates. The proportion of job spells ending in a layoff or firing is 2 pps lower for non-Black than for Black workers.

4.3.3. *Results*

Figure 4 shows the non-parametric hazard estimates using bins of weeks, with no control variables. The patterns are consistent with our model. The one inconsistency, the higher initial point estimate for the layoff hazard for non-Black workers, will be reversed once we add controls and is also

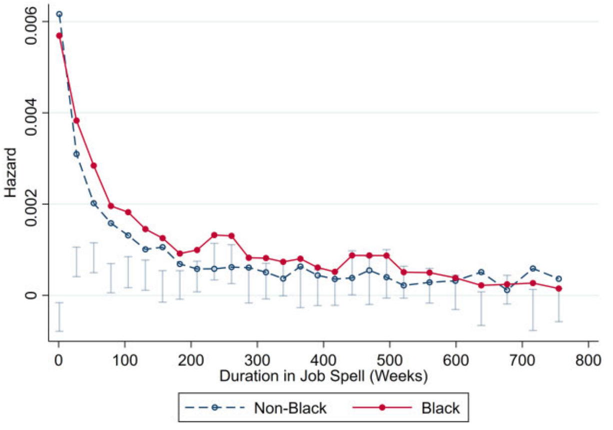

Fig. 4. *Non-Parametric Estimates of the Layoff Hazard by Week Bins, without Control Variables: First Full-Time Spell at Each Employer.*
Notes: We show 90% confidence intervals for the difference between the hazard for Black and non-Black workers. These are based on the estimated SD of the hazard function at the midpoint of the interval, using the fact that the number of failures in the interval is a binomial random variable. If these confidence intervals exclude zero then we can reject the one-tailed test that B − W = 0 with $p = 0.05$. Sample excludes Hispanic workers.

not present in the non-parametric analysis limited to firings. The layoff hazard for Black workers is significantly higher than that for non-Black workers starting in week 26.[47]

We focus our discussion on changes in the absolute gap in the layoff hazard between Black and non-Black workers. If involuntary separations reflect the forces in our model and some race-neutral factor not in our model then the ratio is $(b(t) + c(t))/(w(t) + c(t))$, where $b(t)$ and $w(t)$ are the Black and White rates due to the forces in our model and $c(t)$ is due to some other source. The change in the absolute gap is always $\Delta b - \Delta w$ and thus reflects only the source in our model. Whether the relative gap increases or decreases depends on what is happening to c. When c falls rapidly relative to Δb and Δw, the relative rate will increase. Similarly, the relative rate would increase if both b and w fell by the same amount, but c did not change. Nevertheless, in a robustness check, we present a more parametric model that facilitates examining changes in the relative gap.

The difference in the layoff hazard between Black and non-Black workers starts declining around 1.5 years of tenure, when it falls by roughly 54% relative to the preceding 26-week tenure interval. This decline is statistically significant with $p = 0.05$, using a one-tailed test.[48] The layoff-hazard gap at 1 to 1.5 years of tenure is statistically higher (at the 5% level) than the average of the gaps between 1.5 and 3.5 years.[49] We cannot reject at the 10% level that the gap is zero at around three years of tenure.

[47] The non-parametric survival function for Black workers, constructed using the life-table method, evaluated at the first 26-week period is roughly 86%. While only about 50% of spells are still at risk of failure after this first 26-week period, the decline in spells is largely due to censoring from quits rather than layoffs.

[48] There is a 17-percentage-point decline in the percentage (relative) gap between Black and non-Black hazards between these periods.

[49] We see a roughly 10-percentage-point decline in the percentage gap here.

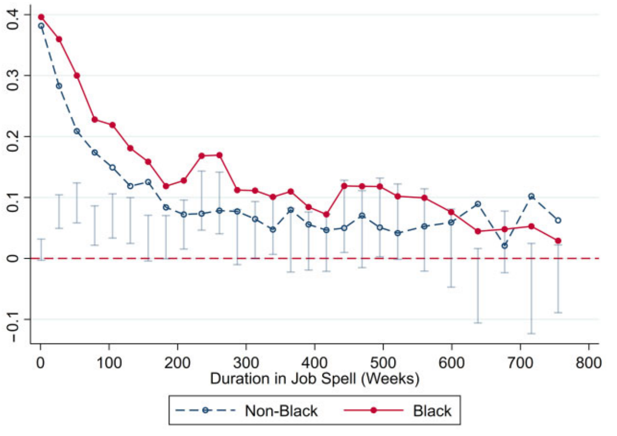

Fig. 5. *Estimates of the Layoff Hazard by Week Bins, Based on a Cox Model Stratified by Race.*
Notes: We show 90% confidence intervals for the difference between the hazard for Black and non-Black workers. Variance of the difference between Black and non-Black hazards is based on a non-parametric bootstrap, and 10,000 bootstrap samples. If these confidence intervals exclude zero then we can reject the one-tailed test that $B - W = 0$ with $p = 0.05$. Sample excludes Hispanic workers.

We can reject a zero gap in some later periods, but they are interspersed among periods in which we cannot reject a zero gap. It is difficult to determine whether these later significant differences are simply spurious. There are many potential tests; all will lead to over-rejection due to multiple hypothesis testing. To provide some discipline to our testing, for each period in which there is an insignificant gap between Black and non-Black workers, we test whether the difference continues to be insignificant when successively adding later periods. We find the interval from 520–60 weeks of tenure is the last in which the hazard difference between Black and non-Black workers was statistically significant on its own or when combined with some number of later periods. Thus, we have no good evidence of significance for anything after 560 weeks of tenure. We conclude that evidence for a notable gap ends somewhere between three and 11 years of tenure.[50]

Online Appendix Figure B3 shows non-parametric plots with hazards by week rather than larger bins, smoothed using kernel smoothers with various bandwidths and kernels, and local linear smoothing. Based on the absence of a gap in the first 26 weeks in Figure 4, it is not surprising that, due to smoothing, most of these plots, which *also do not include control variables*, show a smaller or non-existent gap in the first year followed by an opening of the gap. As before, by year 12, the point estimates suggest no gap in the layoff hazard.

In Figure 5, we present Cox proportional hazard estimates using the multi-week intervals and include the control variables we list in Section 4.3. When adjusting for individual-level covariates, the point estimates suggest a slightly higher, not a lower, hazard for Black workers in the first

[50] Lange (2007) concluded that employers learn about half the information embedded in the AFQT in three years. We analyse the timing of when people are laid off, and show that this is consistent with the result in Lange. We find that 70% of workers would be laid off over 14.5 years of tenure, among workers at risk of being laid off. Roughly half of these would be laid off in the first three years. We calculate these statistics using the default life-table methodology for calculating the survival probability, which assumes that half of the workers who quit over the interval were at risk of being laid off during that period.

26 weeks.[51] Other than this difference, the pattern is similar to that in Figure 4. We again find a statistically significant gap in the layoff hazards from 0.5 to 1.5 years of tenure.

The layoff-hazard gap starts declining at around 1.5 years of tenure, when it falls by roughly 40% relative to the preceding 26-week tenure interval. This decline in the gap is statistically significant at the 10% level, using a one-tailed test.[52] The gap at 1 to 1.5 years of tenure is also statistically higher than the average of the gaps for the four intervals between 1.5 and 3.5 years of tenure at the 10% level and between 1.5 and 4 years of tenure at the 5% level.[53] We cannot reject that the gap is zero at three years of tenure at the 5% level, but can reject at the 10% level using a one-tailed test. Using the test described in the non-parametric analysis above, we have no good evidence that the layoff-hazard gap between Black and non-Black workers is significantly different from zero after 560 weeks of tenure. Our broad conclusion that the gap is statistically negligible starting somewhere between three and 11 years of tenure continues to be valid.

Online Appendix Figure B3 shows additional plots in which we estimate Cox regressions using week rather than larger bins and smoothing hazard contributions using kernel smoothers with various bandwidths and kernels. Similar to Figure 5, these plots show an early gap in the hazards of Black and non-Black workers, which then falls over time. Additionally, the estimated hazard at the earliest tenure is always at least as large for Black workers as for non-Black workers. Online Appendix Figure B3 further shows results using local linear smoothing of the hazard contributions, without any controls. These results are similar to the other non-parametric results in Online Appendix Figure B3.

We also estimated a Cox model with the same covariates, but modelled the percentage gap between the Black and non-Black hazards to be a cubic in seniority. To allow the effect of race on the hazard to vary over time, we include an observation for each job spell at each failure time in the data (as Cox models are only estimated when failures occur). We find that the gap at week one is 3.3%, and this becomes statistically significant at the 10% level at week seven (Online Appendix Figure B4).

Over the range from 1 to 793 weeks, the absolute gap reaches its maximum (over 26-week periods) at 2 to 2.5 years, while the maximum relative gap (64.5%) occurs at 266 weeks (roughly five years). The relative gap then falls, ceasing to be significant at the 5% level at approximately 10.5 years, and reaches zero at roughly 14.5 years.[54] Convergence begins earlier in Figure 4 than in Online Appendix Figure B4 because the latter models the proportional rather than the absolute gap between the hazards. As discussed above, changes in the absolute gap better reflect the forces in our model. In addition, the Online Appendix figure imposes that the ratio is cubic in seniority while our main model is non-parametric in this respect. Together, these results are consistent with our earlier analysis, suggesting large gaps arising by roughly 1 to 1.5 years of tenure and then declining.

4.3.4. *Robustness*

Our main analysis focusses on the hazard of being laid off or fired. Online Appendix Figure B1 shows an alternative specification in which we analyse the hazard only of being fired. The

[51] See Online Appendix Table B2 for coefficients on the covariates. Omitting the region, industry and occupation fixed effects in the Cox estimation yields a larger hazard for non-Black workers than for Black workers in the first 26 weeks.

[52] The percentage gap between Black and non-Black hazards falls over these two periods by roughly 13 percentage points.

[53] Here, too, we see a decline in the percentage gap, of roughly 4 percentage points.

[54] Modelling the percentage gap to be a quartic in seniority yields very similar results.

non-parametric results show that, relative to non-Black workers, Black workers are more likely to be fired starting in the first 26 weeks (though not significantly). This gap becomes significant, starting in weeks 26–52, and then converges. From weeks 638 to 793 (roughly 12.25 to 15.25 years of tenure), no Black workers are fired. Our results from the stratified Cox model are similar.[55]

The incentive to monitor new workers is relevant mainly for individuals hired out of non-employment. For robustness, we identify employment spells for which the individual entered from non-employment and restrict the sample to those spells.[56] We continue to see a gap emerging within the first year of the job spell, and convergence by the twelfth year (Online Appendix Figure B2(c) and (d)). Some readers may feel that the model is more applicable to younger workers. When we restrict the sample to workers no older than 30 at the start of the spell, there is suggestive evidence that the gaps in the first two years are larger in magnitude and remain non-overlapping for an additional 26 weeks (Online Appendix Figure B2(g) and (h)).

When restricting to workers with more than a high school degree, there is much less evidence of differences in hazards between Black and non-Black workers at early tenures (Online Appendix Figure B2(i) and (j)). Similar to our findings above, this suggests that the mechanisms in our model are most relevant for less educated workers.[57]

Online Appendix B.2.6 presents additional analyses showing that our results are robust to treating quits into non-employment as involuntary and treating any gap at an employer as ending the spell.

4.3.5. *Layoff hazard declines with ability for Black workers more than for non-Black workers*

If higher ability reduces the likelihood that a worker makes mistakes or has poor performance, and monitoring allows employers to discover such problems, monitored low-ability workers should be more likely to be fired than monitored workers who are higher ability. This relation should be weaker for non-monitored workers, as mistakes are less likely to be discovered. We test this prediction using AFQT percentile to measure worker ability and estimating a separate Cox model for non-Black and Black workers. We compare the coefficient on AFQT percentile in the two regressions, while recognising the caveats based on the Bond and Lang (2013) critique discussed above.

There is a negative effect of AFQT on the layoff hazard, and the magnitude for Black workers is four times the size of the coefficient for non-Black workers (Table 3; for all coefficients, see Online Appendix Table B3). The difference is significant at the 1% level. We also estimate a Cox model including both Black and non-Black workers, an interaction between Black and AFQT, and between Black and highest grade completed, given the correlation between AFQT and education. This specification yields similar results, and the coefficient on the interaction between Black and AFQT is significant at the 1% level. These results present further evidence consistent with differential monitoring of Black workers.

[55] Based on the Cox estimation, the point estimate of the hazard in the first 26 weeks is only very slightly (0.4%) larger for non-Black workers than for Black workers.

[56] We define individuals as hired out of non-employment if there is more than one week between the end of their last spell and the start of the current spell. Furthermore, we identify individuals as hired from non-employment for their first full-time spell. If the individual did not respond to the following survey after the previous spell, and the year the last spell was reported was at least one year before the current spell was reported (or two years earlier after 1994), the individual is not coded as coming from non-employment.

[57] Online Appendix Figures B2(o)–B2(r) show broadly similar results for men and women.

Table 3. *Differential Effect of AFQT on the Layoff Hazard, by Race.*

	Non-Black	Black	All
AFQT (percentile)	−0.0014	−0.0049	−0.0012
	(0.0007)	(0.0012)	(0.0007)
AFQT × Black			−0.0040
			(0.0013)
Black			−0.188
			(0.222)
Observations	19,293	13,334	32,627

Notes: Conventional SEs are given in parentheses. Coefficients are from a Cox proportional hazard model, using week bin as a unit of time. Each observation is a job spell. We model the layoff hazard, and the failure variable is an indicator for whether the job spell ended because the worker was fired or laid off. The regression additionally includes highest grade completed at spell start (and interacted with Black in column 3), indicators for male, region at spell start (Northeast, North Central, South and omitting West), urban location at spell start, age at spell start, fixed effects for year, occupation (18 groups) and industry (14 groups) all measured at the start of the spell, as well as indicators for whether AFQT, region and urban location are missing. All columns exclude individuals with missing AFQT. Sample excludes Hispanic workers. See the text for details.

5. Conclusion

We develop a model that predicts known disparities between Black and White workers: Black workers earn lower wages, have longer unemployment duration and obtain more education conditional on measured ability. It also predicts one previously unstudied disparity: the layoff hazard is higher for Black workers at low tenure, but the hazard rates converge as tenure increases. In addition, the effect of a measure of unobserved skills on layoffs should be more beneficial for Black than for White workers.

As we have argued previously, while the model, of necessity, relies on some special assumptions, the key elements are (1) that worker productivity is correlated across jobs, (2) that ability is neither perfectly observed or signalled and workers can to some extent hide past firings, (3) that firms therefore use race to statistically infer worker ability, (4) that additional information arrives during employment and is either imperfect, costly or both so that a worker's productivity can never be known perfectly at zero cost and (5) that firms can and do act on new information by firing some workers.

The predictions are largely confirmed. In our stratified Cox models, conditional on observables, Black workers are more likely to be laid off than non-Black workers before one year of seniority. By two years of tenure, this gap has roughly halved. While it is difficult to establish precisely when the layoff-hazard gap falls to 0, we have no evidence of a statistically significant gap after roughly 11 years of tenure. We also confirm that higher unobservable skills, as measured by AFQT, more strongly reduce the likelihood that a Black worker is laid off, relative to a non-Black worker.

Contrary to the model's prediction, our results show that, in the presence of controls, the layoff hazard, while initially larger for Black than for non-Black workers, declines less for Black workers between weeks 1–26 and weeks 26–52 of tenure. Obviously, it is up to the reader to decide how problematic this is. It is plausible that layoffs during the first six months reflect factors not captured by our model and that these initially obscure our model's mechanism. It is not hard to come up with post hoc stories in which some set of White workers would have a higher rate of very short-term employment. As one example, given their better outside options, Whites might be more willing to try jobs where bad matches are readily and quickly observed. The mechanism we underline becomes increasingly important with tenure and dominates after this initial probationary period.

Our message is in some ways depressing. Evidence in this paper and elsewhere suggests that Black workers with high levels of education above the median for American workers can escape the churning equilibrium. However, simply addressing education or human capital disparities between Black and White people need not eliminate labour market disparities. The 'bad equilibrium' in which many Black people find themselves is difficult to escape.

University of Warwick, UK
Boston University, USA, NBER, USA & IZA, Germany
University of Illinois at Urbana-Champaign, USA & IZA, Germany

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Aigner, D.J. and Cain, G.G. (1977). 'Statistical theories of discrimination in labor markets', *Industrial and Labor Relations Review*, vol. 30(2), pp. 175–87.

Autor, D.H. and Scarborough, D. (2008). 'Does job testing harm minority workers? Evidence from retail establishments', *Quarterly Journal of Economics*, vol. 123(1), pp. 219–77.

Bardhi, A., Guo, Y. and Strulovici, B. (2019). 'Spiraling or self-correcting discrimination: A multi-armed bandit approach', Technical Report, Northwestern University.

Becker, G.S. (1971). *The Economics of Discrimination*, 2nd edn., Chicago, IL: Chicago University Press.

Bjerk, D. (2008). 'Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion', Economic Journal, vol. 118(530), pp. 961–82.

Black, D.A. (1995). 'Discrimination in an equilibrium search model', *Journal of Labor Economics*, vol. 13(2), pp. 309–33.

Bond, T.N. and Lang, K. (2013). 'The evolution of the black-white test score gap in grades K–3: The fragility of results', *Review of Economics and Statistics*, vol. 95(5), pp. 1468–79.

Bose, G. and Lang, K. (2017). 'Monitoring for worker quality', *Journal of Labor Economics*, vol. 35(3), pp. 755–85.

Bowlus, A.J. and Eckstein, Z. (2002). 'Discrimination and skill differences in an equilibrium search model', *International Economic Review*, vol. 43(2), pp. 1309–45.

Bowlus, A.J., Kiefer, N.M. and Neumann, G.R. (2001). 'Equilibrium search models and the transition from school to work', *International Economic Review*, vol. 42(2), pp. 317–43.

Bratsberg, B. and Terrell, D. (1998). 'Experience, tenure, and wage growth of young black and white men', *Journal of Human Resources*, vol. 33(3), pp. 658–82.

Bureau of Labor Statistics. (2019). 'National longitudinal survey of youth 1979 cohort, 1979-2010 (rounds 1-24) [computer file]', Columbus, OH: Center for Human Resource Research, The Ohio State University.

Cavounidis, C. and Lang, K. (2015). 'Discrimination and worker evaluation', Working paper, National Bureau of Economic Research.

Charles, K. and Guryan, J. (2008). 'Prejudice and the economics of discrimination', *Journal of Political Economy*, vol. 116(5), pp. 773–809.

Coate, S. and Loury, G.C. (1993). 'Will affirmative-action policies eliminate negative stereotypes?', *American Economic Review*, vol. 83(5), pp. 1220–40.

Coates, T.-N. (2012). 'Fear of a black president', *The Atlantic*, http://www.theatlantic.com/magazine/archive/2012/09/fear-of-a-black-president/309064/# (last accessed: 13 February 2023).

Cornell, B. and Welch, I. (1996). 'Culture, information, and screening discrimination', *Journal of Political Economy*, vol. 104(3), pp. 542–71.

Farber, H. and Gibbons, R. (1996). 'Learning and wage dynamics', *Quarterly Journal of Economics*, vol. 111(6), pp. 1007–48.

Fryer, R.G., Jr. (2007). 'Belief flipping in a dynamic model of statistical discrimination', *Journal of Public Economics*, vol. 91(5–6), pp. 1151–66.

Gehan, E.A. (1969). 'Estimating survival functions from the life table', *Journal of Chronic Diseases*, vol. 21, pp. 629–44.
Guardian. (2014). 'Prince: Black people don't get second chances', *The Guardian*, 25 February 25, http://www.theguardian.com/music/2014/feb/25/prince-black-people-dont-get-second-chances (last accessed: 24 January 2015).
Kalbfleisch, J.D. and Prentice, R.L. (2002). *The Statistical Analysis of Failure Time Data*, 2nd edn., New York: John Wiley.
Kendi, I.X. (2019). *How to Be an Antiracist*, New York, NY: Random House Publishing Group.
Klein, J.P. and Moeschberger, M.L. (2003). *Survival Analysis: Techniques for Censored and Truncated Data*, 2nd edn. New York: Springer.
Lang, K. (1986). 'A language theory of discrimination', *Quarterly Journal of Economics*, vol. 101(2), pp. 363–82.
Lang, K. and Lehmann, J.-Y.K. (2012). 'Racial discrimination in the labor market: Theory and empirics', *Journal of Economic Literature*, vol. 50(4), pp. 959–1006.
Lang, K. and Manove, M. (2003). 'Wage announcements with a continuum of types', *Annales d'Economie et de Statistique*, vol. 71–2, pp. 223–44.
Lang, K., Manove, M. and Dickens, W.T. (2005). 'Racial discrimination in markets with announced wages', *American Economic Review*, vol. 95(4), pp. 1327–40.
Lang, K. and Manove, M. (2011). 'Education and labor market discrimination', *American Economic Review*, vol. 101(4), pp. 1467–96.
Lange, F. (2007). 'The speed of employer learning', *Journal of Labor Economics*, vol. 25(1), pp. 1–35.
Lazear, E.P. (2009). 'Firm-specific human capital: A skill-weights approach', *Journal of Political Economy*, vol. 117(5), pp. 914–40.
Lauermann, S. and Wolinsky, A. (2016). 'Search with adverse selection', *Econometrica*, vol. 84(1), pp. 243–315.
Lundberg, S.J. and Startz, R. (1983). 'Private discrimination and social intervention in competitive labor markets', *American Economic Review*, vol. 73(3), pp. 340–7.
Mabry, M. (2007). *Twice as Good: Condoleezza Rice and Her Path to Power*, New York: Modern Times.
MacLeod, W.B. (2003). 'Optimal contracting with subjective evaluation', *American Economic Review*, vol. 93(1), pp. 216–40.
Masters, A. (2014). 'Statistical discrimination from composition effects in the market for low-skilled workers', *Labour Economics*, vol. 26, pp. 72–80.
Moro, A. and Norman, P. (2004). 'A general equilibrium model of statistical discrimination', *Journal of Economic Theory*, vol. 114(1), pp. 1–30.
National Longitudinal Survey of Youth. (2019). 'NLSY 1979 Appendix 18: Work history data', https://www.nlsinfo.org/content/cohorts/nlsy79/other-documentation/codebook-supplement/nlsy79-appendix-18-work-history-data (last accessed: 15 April 2019).
Neal, D. (1993). 'Supervision and wages across industries', *Review of Economics and Statistics*, vol. 75(3), pp. 409–17.
Oyer, P. and Schaefer, S. (2000). 'Layoffs and litigation', *RAND Journal of Economics*, vol. 31(2), pp. 345–58.
Panel Study of Income Dynamics. (2020). 'Public use dataset', Ann Arbor, MI: Survey Research Center, Institute for Social Research, University of Michigan, https://simba.isr.umich.edu/data/data.aspx (last accessed: 16 January 2023).
Peski, M. and Szentes, B. (2013). 'Spontaneous discrimination', *American Economic Review*, vol. 103(6), pp. 2412–36.
Postel-Vinay, F. and Robin, J.-M. (2002). 'Equilibrium wage dispersion with worker and employer heterogeneity', *Econometrica*, vol. 70(6), pp. 2295–360.
Reid, A. (2015). 'Black head coaches don't get second chances', *St. Louis American*, 31 August, https://www.stlamerican.com/sports/sports_columnists/sports_eye/black-head-coaches-don-t-get-second-chances/article_00d64010-4c4b-11e5-9045-57588bb5b35c.html (last accessed: 13 October 2023).
Rosen, A. (1997). 'An equilibrium search-matching model of discrimination', *European Economic Review*, vol. 41(8), pp. 1589–613.
Rubinstein, A. (1982). 'Perfect equilibrium in a bargaining model', *Econometrica*, vol. 50(1), pp. 97–109.
Ruggles, S., Flood, S., Goeken, R., Schouweiler, M. and Sobek, M. (2022). 'IPUMS USA: Version 12.0 [dataset]', Minneapolis, MN: IPUMS, https://doi.org/10.18128/D010.V12.0.
Spencer, D. (2014). 'For black youth in America, there are no second chances', *emPower*, 5 February, https://web.archive.org/web/20151115010821/ (last accessed: 16 October 2023).
Wozniak, A. (2015). 'Discrimination and the effects of drug testing on black employment', *Review of Economics and Statistics*, vol. 97(3), pp. 548–66.

The Economic Journal, **134** (*February*), 485–514 https://doi.org/10.1093/ej/uead079
Advance Access Publication Date: 27 September 2023