Antidiscrimination Enforcement and the Problem of Patronization

By Stephen Coate and Glenn Loury*

A government determined to prohibit race or sex discrimination in employment faces a difficult enforcement problem. With unlimited information a regulator might, by observing the outcome of every employment decision made by a firm, determine whether that firm is using the same criteria to select among applicants from different groups. However, such a wealth of information is not available in practice. Even if it were, the sheer scale of the labor market would make it impossible to monitor effectively every transaction. Therefore, enforcement of antidiscrimination laws must rest upon means other than the exhaustive, direct observation of individual employment decisions.

One response is to compare the representation of minorities and women in a particular firm with their representation in the population. To the extent that there is a significant divergence, the firm could be asked to account for it. Absent a compelling justification, the difference would be presumed evidence of bias. This technique is implicit in one of the key provisions of the Civil Rights Act of 1991, which holds any employment practice having a "disparate impact" on women or minorities to be unlawful unless the firm can demonstrate that the practice constitutes a "business necessity" (Richard Epstein, 1992).

However, objections can be raised against this type of "statistical" enforcement policy. If the distribution of skills differs across groups, the representation of women and minorities in the population may overstate their actual employment rates under neutral hiring rules. The burden then falls to firms to prove that any difference is due to a disparity in group skills. This may be difficult to establish to the satisfaction of a court, and anticipating this difficulty, firms may respond to this enforcement regime by adopting hiring rules biased *in favor* of women and minorities.

This concern about "quotas" motivated much of the conservative criticism of the Civil Rights Act of 1991. Such criticism misses a basic point, however: one consequence of employment discrimination is that it harms the incentives for workers to acquire skills. If workers expect to face biased hiring rules, their anticipated returns from investing in job-relevant skills are reduced. Thus, group disparities in skills may simply reflect the presence of employment discrimination.

Hence, a regulator may be justified in placing a "burden of proof" on firms whose work force exhibits substantial underrepresentation of women or minorities in certain jobs. The ultimate effect of this burden may simply be to induce firms that would otherwise discriminate to offer equal opportunities to all workers and to eliminate skill disparities across groups. This is the view taken by many liberal advocates of stronger civil-rights laws. Yet, this view also overlooks an important point: statistical enforcement policies will not necessarily move workers' incentives in the right direction (Coate and Loury, 1992). If the policy forces firms to "patronize" some workers by setting lower standards for them, then the workers may be persuaded that they can get desired jobs without making costly investments in skills. However, if fewer members of some group acquire skills, firms will be forced to continue patronizing them in or-

*Coate: Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104-6297; Loury: Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215.

der to achieve parity. Thus, skill disparities might persist, or even worsen, under such policies.

This paper develops a simple taste-discrimination model which illustrates these problems in the enforcement of antidiscrimination laws. In the model, discrimination reduces the incentives to invest and hence creates skill disparities between groups. These disparities are not necessarily eliminated by statistical enforcement policies because of the adverse incentives created by patronization. The model suggests that a gradual policy in which representation targets are ratcheted up through time will be more likely to eliminate both discrimination and skill disparities than a radical policy that demands immediate proportional representation.

I. Discrimination and the Acquisition of Skills

Consider a model of the labor market in which each of a large number of firms hires its work force from a common population. A large number of prospective employees (drawn randomly from this population) approach each firm seeking employment. These workers belong to one of two identifiable groups, denoted B or W; λ is the fraction of W's in the population and hence is also the probability that a given applicant for employment is a W. When a firm encounters a worker it must decide whether to "accept" or "reject" the applicant. A firm engages in discrimination if it uses a different rule when deciding whether to accept B's than it does when deciding about W's.

Individual workers are either "qualified" or "unqualified," and firms can observe a worker's qualifications before deciding whether to accept him. Acceptance yields a worker the gross payoff ω irrespective of his qualifications, and a rejected worker's gross payoff is normalized to zero. The direct monetary return to a firm from accepting a worker is $x_q > 0$ if he is qualified, and $-x_u < 0$ if he is unqualified.

Firms are taste discriminators in that they experience some psychic cost from accepting a B (Gary Becker, 1957). This cost is greater, the larger the ratio of B's to W's among the pool of acceptees. Specifically, for some $\gamma > 0$, if an accepted worker is a B, then a firm's return, net of psychic cost, is $x_q - \gamma r$ if he is qualified and $-x_u - \gamma r$ if he is unqualified, where r is the ratio of accepted B's to W's for that firm. The payoff to any firm from rejecting any worker is taken to be zero. To keep things simple, the payoff parameters x_q, x_u, ω and γ are taken to be exogenous. In particular, firms are not allowed to offer lower wages to B's as compensation for their psychic costs. Thus the focus is on the implications of discriminatory hiring in a world where equal pay laws are perfectly enforced.

To become qualified, workers must make some costly *ex ante* investment. This investment may be thought of either as acquiring knowledge or life skills. The cost of becoming qualified varies in the population but is distributed in the same way among B's and W's. Let c denote an individual's investment cost, and let $G(c)$ be the fraction of workers in either group with cost no greater than c. Workers decide, prior to encountering a firm, whether to make this investment, acquiring the skill if the expected payoff from doing so is no less than its costs. We assume that G is smooth and increasing with $G(0)=0$ and $G(\omega)<1$. This last assumption guarantees that not all workers will invest, even when being qualified is both necessary and sufficient for acceptance.

The timing of the interaction between workers and firms is as follows. First, the various workers decide whether to invest. They are then randomly matched with firms. Finally, each firm, upon observing the group identities and qualifications of its applicants, makes its acceptance decisions, and payoffs are received. To represent the agents' behavior formally, let $I(i,c)$, $i=\mathrm{B}$ or W and $c \geq 0$, be a strategy for workers, giving the probability of investing as a function of group identity and cost. Let $A(i,k)$, $k = \mathrm{u}$ or q, be a strategy for firms, giving the probability of accepting a worker as a function of his group identity and state of qualification. Then an *equilibrium* is a pair of strategies (I, A) such that each is a best response to the other. We shall assume that firms' taste for discrimination is sufficiently "strong" in the following sense:

ASSUMPTION 1:

$$\gamma > \max\{\lambda x_q / 2(1-\lambda);\; x_q^2 / 4x_u\}.$$

To find the equilibrium consider first the behavior of firms. No firm accepts an unqualified B or rejects a qualified W. Doing so would both lower a firm's monetary returns and increase its psychic costs. Thus, $A(\mathrm{B}, \mathrm{u}) = 0$ and $A(\mathrm{W}, \mathrm{q}) = 1$. We also claim that firms reject unqualified W's, and accept qualified B's if their ratio among the accepted applicants is not too great. To see this, suppose a firm follows an employment rule which results in the acceptance of z_b B's and z_w W's, leaving a ratio of B's to W's among its accepted workers of $r \equiv z_b / z_w$. This firm incurs the psychic cost $\gamma r z_b$ due to its taste for discrimination. Thus, its marginal psychic cost of accepting another B is $2\gamma r$, while its marginal psychic benefit of accepting another W is γr^2. Since the monetary benefit of accepting a qualified B is x_q, it follows that a maximizing firm accepts another qualified B so long as $x_q \geq 2\gamma r$. That is, a firm accepts qualified B's so long as the B/W ratio among the accepted applicants satisfies $r \leq r^*$, where $r^* \equiv x_q / 2\gamma$.

Furthermore, since the monetary cost of accepting an unqualified *W* is x_u, it pays for a firm to do so only if $x_u \leq \gamma r^2$. We have just shown that firms never allow r to exceed r^*; so accepting an unqualified W will never pay if $x_u > \gamma (r^*)^2$, a condition which is guaranteed by Assumption 1. Thus, a firm behaving optimally will reject all unqualified workers, accept all qualified W's, and accept a qualified B if and only if $r \leq r^*$.

With the workers' strategy $I(i, c)$ given, let π_i be the fraction of group i who invest: $\pi_i \equiv \int I(i, c)\, dG(c)$. Because workers are matched with firms randomly and in large numbers, the shares of qualified W's and B's in a firm's applicant pool are (approximately) $\lambda \pi_w$ and $(1-\lambda)\pi_b$, respectively. Let $\bar{r} = (1-\lambda)/\lambda$ be the ratio of B's to W's in the population. By the foregoing reasoning, if $\bar{r}(\pi_b / \pi_w) \leq r^*$ a firm's best response involves accepting all of its qualified B applicants, while if $\bar{r}(\pi_b / \pi_w) > r^*$ a firm will accept some, but not all, qualified B's. This means that the typical qualified B is accepted with probability $A(\mathrm{B}, \mathrm{q}) = \delta(\pi_b, \pi_w)$, where

$$(1) \quad \delta(\pi_b, \pi_w) \equiv \min\{(\pi_w r^* / \pi_b \bar{r}); 1\}.$$

Thus, firms will discriminate against some B's if workers qualify at rates such that $\pi_b / \pi_w > r^* / \bar{r}$.

Now consider workers' best response to this behavior by firms. Since W's are accepted if and only if they invest, their return from investing is ω; so those W's with $c \leq \omega$ invest, implying $\pi_w = G(\omega)$. Since B's are rejected if they do not invest and face some probability δ of being accepted if they do, their return from investing is $\omega\delta$; so those B's with $c \leq \omega\delta$ invest, implying $\pi_b = G(\omega\delta)$. We conclude from this discussion that in equilibrium the acceptance probability for qualified B's, δ^*, must solve the following equation:

$$(2) \quad \delta = \min\{G(\omega) r^* / G(\delta\omega)\bar{r}; 1\}.$$

Since Assumption 1 implies $r^* < \bar{r}$, it is easily seen that (2) has a unique solution δ^*, with $0 < \delta^* < 1$. That is, equilibrium must involve discrimination against B's.

We conclude, therefore, that when the taste for discrimination is strong enough the only possible outcome is that firms accept qualified W's, reject unqualified B's and W's, and accept qualified B's with some probability $\delta^* < 1$. This discrimination against B's lowers their incentive to invest. As a result, W's become qualified at a higher rate than B's. Hence, B's are underrepresented among the accepted, in terms of both their presence in the population and, to a lesser extent, their presence among the skilled $(r^* / \bar{r} < \delta^* < 1)$.

II. The Impact of Antidiscrimination Enforcement

As mentioned above, a government intent on fighting discrimination here might begin by insisting that firms accept all qualified workers irrespective of group identity. Enforcing this mandate on a case-by-case basis would require knowledge of the qualifications of the applicant and the acceptance decision of the firm in each and every case. Moreover, it would also be necessary to verify this information in court. Given the practical difficulties of this approach, the use of a statistical enforcement strategy, based on observed aggregate outcomes at the firm level, may be quite attractive.

We can depict such an enforcement approach in our model as follows. Suppose that the regulator selects some target ratio $\hat{r}$ and announces that any firm found with a ratio of B's to W's less than $\hat{r}$ will face costly legal proceedings. Let the anticipated costs of these proceedings be so great that no firm wants to risk being found in violation. Then firms will adapt to this regulatory regime by altering their acceptance rules to ensure that $r \geq \hat{r}$. We will assess the effects of this kind of regulation by analyzing how the equilibrium outcome of our simple model changes under this constraint.

The objective of full proportional representation would correspond to a target level $\hat{r} = \bar{r}$. However, for reasons discussed later, a regulator might also want to consider more modest objectives. Thus we will conduct the analysis under the assumption that $\hat{r}$ lies somewhere between r^* and $\bar{r}$: $\hat{r} \in (r^*, \bar{r}]$. That is, the target is set to increase the representation of B's, but not beyond their relative numbers in the population.

To analyze the impact of this constraint, notice first that if the taste for discrimination (γ) is large enough, the enforcement policy may cause the market to collapse, with no workers being accepted and none investing. To see this, observe that the constraint must be binding for all firms, since a firm in strict compliance (with $r > \hat{r}$) can reject another B, thereby reducing its psychic costs by $2\gamma r$ while incurring a monetary loss of at most x_q (should the rejected B be qualified). This rejection pays since $r > \hat{r} > r^*$. Furthermore, since the constraint is binding, firms cannot gain by accepting unqualified W's. Thus, the firm's problem under the regulatory constraint may be reduced to choosing how many qualified W's to accept, with B's then accepted at a rate that ensures compliance.

Since compliance requires there to be $\hat{r}$ B's for each W among the accepted, the net benefit from accepting a qualified W, taking account of monetary and psychic returns and of the regulatory constraint is $x_q + \hat{r}x_q - \gamma\hat{r}^2$ if the marginal B accepted is qualified, and it is $x_q - \hat{r}x_u - \gamma\hat{r}^2$ if the marginal B is unqualified. Thus, if $x_q + \hat{r}x_q - \gamma\hat{r}^2 < 0$ firms will reject all their applicants, and the collapse of the market is assured. To avoid this outcome for all $\hat{r} \in (r^*, \bar{r}]$ requires that $x_q + \bar{r}x_q - \gamma\bar{r}^2 > 0$, which amounts to the following:

ASSUMPTION 2:

$$\gamma < \left[x_q/(1-\lambda)\right]\left[\lambda/(1-\lambda)\right].$$

Assumptions 1 and 2 together simply state that firms dislike B's enough to discriminate against them in the absence of regulation, but not so much as to forgo operation altogether if required to employ qualified B's at a rate equal to their presence in the population.

Under Assumption 2, firms gain by accepting qualified W's so long as compliance with the regulation can be maintained by accepting qualified B's. If the ratio of B's to W's among a firm's qualified applicants is less than $\hat{r}$, the firm will need to consider whether it pays to accept any unqualified B's. From the above discussion, this will be worthwhile if

$$x_q > x_u\hat{r} + \gamma\hat{r}^2. \tag{3}$$

Hence, there are two cases of interest, depending on whether or not (3) holds. If (3) fails (case 1), firms accept the maximal number of qualified W's consistent with being able to remain in compliance by accepting only qualified B's. If (3) holds (case 2), firms accept all qualified W's, and as many B's as necessary to remain in compliance. For $\hat{r}$ near $\bar{r}$, (3) is more demanding than Assumption 2. Nevertheless, it can be shown that values for the parameters x_q, x_u, and γ exist satisfying (3) for all $\hat{r} \in (r^*, \bar{r}]$, and satisfying Assumption 1, as long as $\lambda > 1/2$.

It is now possible to describe how the government's statistical enforcement policy alters the equilibrium of the model.

PROPOSITION 1: *In case* 1 *there is a unique equilibrium for each statistical enforcement policy* $\hat{r} \in (r^*, \bar{r}]$ *in which all unqualified workers are rejected, all qualified W's are accepted, and qualified B's are accepted with probability* $\delta(\hat{r})$, $0 < \delta(\hat{r}) \leq 1$. *This probability* $\delta(\hat{r})$ *exceeds the laissez-faire acceptance rate* δ^* *and is strictly increasing in* $\hat{r}$; $\delta(\bar{r}) = 1$.

PROOF:

If B's and W's invest at rates π_b and π_w, then the fractions of a firm's applicants who are qualified B's and W's are $(1-\lambda)\pi_b$ and $\lambda\pi_w$, respectively. In case 1, firms accept none of the unqualified applicants; among the qualified applicants, firms accept all W's and some B's if $\hat{r}/\bar{r} < \pi_b/\pi_w$; and all B's and some W's are accepted if $\hat{r}/\bar{r} > \pi_b/\pi_w$. But when $\hat{r} \in (r^*, \bar{r}]$ an equilibrium with $\hat{r}/\bar{r} > \pi_b/\pi_w$ is impossible, since accepting all B's and not all W's from among the qualified means that $\pi_b > \pi_w$, and so $\hat{r}/\bar{r} > 1$, a contradiction. Thus, equilibrium necessarily entails acceptance from among the qualified of all W's, and some fraction $\delta = (\hat{r}/\bar{r})(\pi_w/\pi_b)$ of B's, where $\pi_w = G(\omega)$ and $\pi_b = G(\delta\omega)$. Thus, the equilibrium acceptance rate for qualified B's, $\delta(\hat{r})$, solves: $\delta G(\delta\omega) = (\hat{r}/\bar{r})G(\omega)$. This solution is unique, exceeds δ^* for all $\hat{r} \in (r^*, \bar{r}]$, is increasing in $\hat{r}$, and equals 1 when $\hat{r} = \bar{r}$.

In case 1, therefore, the use of a statistical enforcement strategy is sure to produce desirable results. Setting the target equal to the population ratio implies an outcome with no discrimination and no skill disparity between groups. More generally, a stricter target leads to less discrimination by the firm and less of a skill disparity between the groups. Unfortunately, matters are not so comforting in case 2.

PROPOSITION 2: *In case* 2, *the equilibrium described in Proposition* 1 *continues to exist. In addition, however, "patronizing" equilibria may also exist. In these equilibria all qualified workers are accepted, all unqualified W's are rejected, but unqualified B's are accepted with positive probability. A necessary and sufficient condition for the existence of a patronizing equilibrium is that, for some* $\delta \in (0,1)$, $\delta[1-G(\omega\delta)] \geq [1-(\hat{r}/\bar{r})G(\omega)]$.

PROOF:

In case 2, firms accept unqualified B's if and only if this is necessary to remain in compliance when accepting all the qualified W's. In the equilibrium of Proposition 1 compliance can be achieved by accepting all qualified W's and no unqualified B's, since $\pi_b/\pi_w \geq \hat{r}/\bar{r}$; so this remains an equilibrium in case 2. We seek to identify other (patronizing) equilibria, where $\pi_b/\pi_w < \hat{r}/\bar{r}$, in case 2. In such an equilibrium, to comply while accepting all qualified W's, firms must accept all qualified B's and the fraction $\sigma = [(\hat{r}/\bar{r})\pi_w - \pi_b]/(1-\pi_b)$ of unqualified B's. However, then, the return to a B from investing is $\delta = 1 - \sigma = [1-(\hat{r}/\bar{r})\pi_w]/(1-\pi_b)$. Hence the fractions $\pi_b = G(\omega\delta)$ and $\pi_w = G(\omega)$ of B's and W's would invest. Therefore, a patronizing equilibrium exists if $\delta = [1-(\hat{r}/\bar{r})G(\omega)]/[1-G(\omega\delta)]$ has a solution for some $\delta \in (0,1)$. A necessary and sufficient condition for this to occur is that $\delta \geq [1-(\hat{r}/\bar{r})G(\omega)]/[1-G(\omega\delta)]$ for some $\delta \in (0,1)$. Indeed, if this inequality holds strictly, then at least two patronizing equilibria exist.

Therefore, in case 2 and under the above stated condition, the use of a statistical enforcement policy can lead firms to patronize B's in equilibrium. Because they are willing to accept unqualified B's in order to meet the hiring target, firms may act so as to lower the incentive for B's to invest, thereby inducing a skill disparity. Indeed, the skill disparity may actually widen as a result of the regulator's intervention. Notice that the condition in Proposition 2 will be more easily satisfied when $\hat{r}$ is larger. That is, patronizing equilibria are more likely to exist when the target is more ambitious.

To get an intuitive grasp of what is going on in these patronizing equilibria, consider how firms might be expected to react to the initial imposition of some hiring target $\hat{r} > r^*$. Prior to the introduction of the regulation there is a "surplus" of qualified B's, in that more B's invest than find employment. Therefore, when the target is modest, firms anticipate meeting it by accepting more qualified B's without having to take on any of the unqualified. This response increases the incentives for B's to invest, and the new equilibrium is as described in Proposition 1, with more B's employed and the skills gap narrowed.

However, if case 2 applies and if the target is sufficiently ambitious, then firms perceive a "shortage" of qualified B's, relative to the numbers needed to be in compliance, while accepting all qualified W's. They therefore switch from discriminating against qualified B's to discriminating in favor of unqualified B's. This response can lower incentives for B's to become skilled, leading to a patronizing equilibrium such as that described in Proposition 2.

How likely is it that such an equilibrium would arise? We address this question by imagining successive cohorts of workers interacting with firms over time. A dynamic adjustment process is defined for a given hiring target $\hat{r}$ by assuming that firms hire optimally from the workers in each cohort, while investment decisions in cohort $t+1$ are based on the acceptance rules applied by firms to cohort t. By iterating this process we trace out a long-run response to the enforcement policy. The fraction of B's investing initially is denoted by π_b^0. Since, in case 2, firms always accept all of the qualified and none of the unqualified W's, we know that the constant fraction $\pi_w = G(\omega)$ of W's invests in each cohort.

Let π_b^t be the investment rate of B's in cohort t. If $\pi_b^t/\pi_w \geq \hat{r}/\bar{r}$, firms accept a qualified B with probability $\delta^t \equiv (\hat{r}/\bar{r})(\pi_w/\pi_b^t)$ to comply, so B's in cohort $t+1$ invest at rate $\pi_b^{t+1} = G(\omega\delta^t)$. If $\pi_b^t/\pi_w < \hat{r}/\bar{r}$, all qualified B's and the fraction $\sigma^t = [(\hat{r}/\bar{r})\pi_w - \pi_b^t]/(1-\pi_b^t)$ of unqualified B's are accepted, so B's in cohort $t+1$ invest at rate $\pi_b^{t+1} = G(\omega(1-\sigma^t))$. Thus $\{\pi_b^t\}$ solves:

$$(4) \quad \pi_b^{t+1} = G\big(\omega[\hat{r}/\bar{r}][\pi_w/\pi_b^t]\big) \qquad \text{for } \pi_b^t/\pi_w \geq \hat{r}/\bar{r}$$

$$\pi_b^{t+1} = G\big(\omega[1-(\hat{r}/\bar{r})\pi_w]/[1-\pi_b^t]\big) \qquad \text{for } \pi_b^t/\pi_w < \hat{r}/\bar{r}.$$

The stationary points of this difference equation correspond exactly to the fraction of B's who invest in the equilibria of our model. Hence, if a patronizing equilibrium exists then (4) has a stationary point π_b^* at which $\pi_b^*/\pi_w < \hat{r}/\bar{r}$. Detailed study of (4) leads to the following result.

PROPOSITION 3: *Given any enforcement policy $\hat{r} \in (r^*, \bar{r}]$, if a patronizing equilibrium exists under $\hat{r}$ then there is a critical value $\tilde{\pi}_b(\hat{r}) \in (0, \pi_w]$, increasing in $\hat{r}$, such that* (4) *converges to a patronizing equilibrium if $\pi_b^0 \leq \tilde{\pi}_b(\hat{r})$, while it converges to the equilibrium described in Proposition 1 if $\pi_b^0 > \tilde{\pi}_b(\hat{r})$.*[1] *Furthermore, if $\omega G'(\omega)/[1-G(\omega)] > 1$ (true for uniform or exponential cost distributions when ω exceeds the average investment cost) then a patronizing equilibrium always exists for $\hat{r}$ close enough to $\bar{r}$, and moreover, $\tilde{\pi}_b(\bar{r}) = \pi_w$.*

(The proof is available from the authors upon request.)

This result has an important implication. Note that $\pi_b^0 < \pi_w$ necessarily, for otherwise there is no need for regulation. Hence,

[1]This discussion assumes that the equilibrium of Proposition 1 is locally stable under the dynamic adjustment process specified in (4). Conditions on G sufficient to assure this are easily deduced. Notice that when multiple patronizing equilibria exist, at least one of them must be locally stable under the adjustment process (4).

given that inequality (3) holds when $\hat{r} = \bar{r}$ and that $\omega G'(\omega)/[1 - G(\omega)] > 1$, Proposition 3 implies that *a patronizing outcome is the inevitable result of the regulator seeking fully proportional representation* [$\hat{r} = \bar{r}$]. However, since the threshold investment rate $\tilde{\pi}_b(\hat{r})$ is increasing in $\hat{r}$, the regulator could avoid an equilibrium in which B's are patronized by setting a more modest target. One can show, for example, that a target equal to the laissez-faire ratio of qualified B's to W's [$\hat{r} = \bar{r}G(\delta^* \omega)/G(\omega)$] never leads to patronization.

These results suggest that antidiscrimination enforcement embodies an awkward trade-off: a policy of proportional representation of B's risks inducing a patronizing outcome, while a more modest target, though less prone to that problem, will not fully eliminate discrimination. Yet our model also suggests a way around this trade-off. Rather than settling immediately upon the proportional target $\hat{r} = \bar{r}$, the regulator could instead operate a gradual policy, with the target being ratcheted up in a series of steps. Suppose the adjustment process described in (4) operates quickly, relative to the rate at which the policy is changed. Then, if the regulator always sets the new target r' so that the currently prevailing investment rate among B's, π_b, satisfies $\pi_b > \tilde{\pi}_b(r')$, a patronizing outcome can be avoided. However, by Proposition 1, each time the target is raised the rate of investment among B's improves, which permits the target to be raised further without risk of patronization. By proceeding in this way, this process will eventually eliminate both discrimination and also skill disparities between the groups.

REFERENCES

Becker, Gary, *The Economics of Discrimination*, Chicago: University of Chicago Press, 1957.

Coate, Stephen and Loury, Glenn, "Will Affirmative Action Policies Eliminate Negative Stereotypes?" Department of Economics Working Paper No. 3, Boston University, June 1992; *American Economic Review* (forthcoming).

Epstein, Richard, *Forbidden Grounds: The Case Against Employment Discrimination Laws*, Cambridge, MA: Harvard University Press, 1992.