ECONOMIC RESEARCH REPORTS

COLOR-BLIND IS NOT COLOR-NEUTRAL: DISADVANTAGE AND AFFIRMATIVE ACTION

BY

Jean-Pierre Benoit

RR # 94-27 September 1994

C. V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003

Color-Blind is not Color-Neutral: Disadvantage and Affirmative Action

by

Jean-Pierre Benoît*
Economics Department
and School of Law
New York University

June 1994
Revised September 1994

*This work was begun while the author was visiting at the California Institute of Technology and the Univeristy of Southern California Law school. He thanks Orley Ashenfelter, Lewis Kornhauser, Vijay Krishna, Enza Orlandi, Glenn Loury, Scott Page, James Strnad, and Quang Vuong for their comments, and the CV Starr Center for its support.

Abstract

Employers or universities determine the qualification of applicants based upon the results of a test. Members of socio-economically disadvantaged groups tend to score less well than equally qualified members of other groups. As a result, color-blind practices discriminate against disadvantaged groups. This discrimination may persist even if rational firms realize that the test is biased. An affirmative action program may be needed to achieve color-neutral results.

Jean-Pierre Benoît

Color-Blind is not Color-Neutral: Disadvantage and Affirmative Action

Though the color of money may be green, the color of those who earn it is not so uniform. United States figures for 1992 show a median weekly income of $518 for white males, $380 for black males and $345 for Hispanic males[1]. While %10.1 of employed workers were African-American, only 6.5% of managerial and professional workers were African-American; corresponding figures for Hispanics were 7.6% and 3.9%. On the other hand 17.5 % of those in service jobs were African-American and 11.3 % were Hispanic. Unemployment rates for whites were %6.5, %14.1 for blacks, and 11.4% for Hispanics. Gender differences are also pronounced. For instance, the overall male to female median wage ratio for 1992 was 1.32.

While these disparities are striking, little consensus has been reached on, what, if anything should be done about them. One of the more controversial proposals has been affirmative action. But although affirmative action has provoked much discussion, there has not been very much formal analysis. A recent exception is a paper by Coate and Loury (1993)[2]. They argue, somewhat surprisingly, that even if group disparities can be attributed to discrimination on the part of firms, affirmative action may not be an appropriate remedy. This paper follows the basic framework of Coate and Loury closely, but reaches quite different conclusions. Indeed, I show that even if workers are not being discriminated against by firms, affirmative action may be needed to correct socio-economic inequities.

The outline of the model used is as follows. Firms are interested in hiring or promoting qualified workers. In order to determine whether or not a worker is qualified, the firms administer an imperfect test; while higher scores tend to indicate that a worker is qualified, it is possible both that a qualified worker receives a low score and that an unqualified worker receives a high score. Workers fall into two identifiable groups, B's and W's. Due to socio-economic inequities (for instance B's attend poorer high schools) B's tend to score lower on the test than do W's. That is, given a B worker and a W worker equally qualified, or unqualified, for the job at hand, the B worker has a lower expected test score. It is unsurprising then, that if firms set a single standard for all workers, B's may end up underrepresented in good jobs. Thus, color-blind employment practices discriminate against the disadvantaged B workers. Furthermore, B's are particularly disfavored in high skill jobs.

It may seem, however, that a profit maximizing firm, realizing that the lower scores of B's are an artifact of social conditions, will lower the test standards for B's in order to hire the equally qualified though lower scoring B's. In this way, B's and W's will end up equally represented and there should be no need for an affirmative action program. This is not the case. If B's start out sufficiently disadvantaged, then absent some kind of government

[1]All data is taken from the Statistical Abstract of the United States(1993) . The Abstract provides information for these three racial groups

[2]Other papers include Welch(1976). See Coate and Loury also for a fuller survey of related literature, including Arrow(1973) whose model bears some relation.

program, discriminatory results will persist. The reason is that the initial standard discourages many B's from becoming qualified. Given the low proportion of qualified B's, it is no longer rational for the individual firms to reduce standards for them. However, if the firms as a group were to lower standards sufficiently, B's would be induced to become qualified at the same rate as W's. This can be accomplished through an affirmative action program.

Section 2 reviews the Coate and Loury model. Section 3 offers a criticism of their conclusion, while maintaining their assumptions. Section 4 introduces the assumption of unequal test performane and is the main contribution of this paper. Section 5 contrasts the assumptions of section 4 with an assumption of Coate and Loury. Sections 6 and 7 provide final remarks.

2. The Coate and Loury Model

In this section I present Coate and Loury's basic model and their main conclusion. My presentation is fairly detailed since their model forms the basis of the subsequent analysis[3].

There is a large (competitive) number of identical firms and a much larger number of workers. Each firm has two job levels, skilled and unskilled (the unskilled job can also be interpreted as not hiring the worker). A worker is either qualified for the skilled job or unqualified. All workers can perform the unskilled job satisfactorily. Workers prefer the skilled job to the unskilled job, while firms want to assign only qualified workers to the skilled job. A worker's gross benefit from assignment to the skilled job is w (equal pay for equal work laws are in effect, so that a worker's pay depends only on his job assignment). An employer's net return to assigning a qualified worker to the skilled job is $x_q>0$ and her net return to assigning an unqualified worker to a skilled job is $-x_u < 0$. The worker's return to being assigned to an unskilled job and the employer's benefit from assigning any worker to an unskilled job are both normalized to 0.

In order to determine whether or not a worker is qualified, the firm administers a test, or evaluates the worker's initial job performance. Although the word "test" is used the evaluation procedure may be informal and may not be easily understood by someone outside the firm. The firm receives a signal $\theta \in [0,1]$ which is subject to type I and type II errors. That is, it is possible, though relatively unlikely, that a qualified worker scores low on the test, and that an unqualified worker scores high on the test. The density function of test scores for qualified and unqualified workers is given by $f_q(\theta)$ and $f_u(\theta)$. $f_u(\theta)/f_q(\theta)$ is assumed to be continuous, strictly decreasing, and positive and finite on [0, 1], which (more than) implies that $F_q(\theta) < F_u(\theta)$ -- qualified workers tend to score higher than unqualified workers.

Workers fall into two identifiable groups B and W. Workers in different groups will

3 There are some minor differences in the way their model is presented here, but I follow them closely, even using some of their descriptions verbatim. Also their paper contains several variations and discussions (including arguments for affirmative action and other remedies) that I do not treat. The reader is encouraged to consult their original paper.

sometimes be referred to as being of different "colors", although the workers could just as well be grouped by gender or any other distinguishing characteristic. Ex ante these groups are identical. All workers start out life unqualified and can become qualified by making an investment decision (such as working hard at school, or studying their job manuals at home). The cost of making this investment varies across worker. Let G(c) be the fraction of workers with investment cost at most c (this function is the same for both groups of workers). G is assumed to be continuous with positive density over the relevant range and $G(0) = 0$.

The test is administered after the workers' investment decisions. Following the test the worker is randomly assigned to a firm. Alternatively, we can say that a firm does not particularly favor its own workers over others. Given a fraction π of qualified workers in a population, the posterior probability that a worker is qualified given that he has emitted a test signal θ is given by

$$p(\pi, \theta) = \pi f_q(\theta)/[\pi f_q(\theta) + (1-\pi)f_u(\theta)]$$

The firm optimally assigns the worker to the skilled job if:

$$p(\pi,\theta)x_q - (1-p(\pi, \theta)x_u \geq 0.$$

This yields a threshold rule for the employer: assign a worker to a qualified task if and only if he emits a signal no less than the standard $s^*(\pi)$, where

$$s^*(\pi) \equiv \min\{\theta\epsilon[0,1]: x_q/x_u \geq [(1-\pi)/\pi]f_u(\theta)/f_q(\theta)\} \text{ or 0 if the inequality never holds.}$$

For a worker the benefit ß(s) to investing is given by the increase in the probability of being assigned to a skilled job times the benefit from such an assignment:

$$ß(s) \equiv w[F_u(s) - F_q(s)]$$

A worker with cost c invests if and only if $c \leq ß(s)$. Therefore, if the standard is s, the fraction of workers investing, i.e., the fraction of qualified workers, is given by G(ß(s)).

The situation is depicted in figure 1[4]. Given a signal s, the employer's line EE(s) gives the fraction π of the population that would have to be qualified in order for it to be optimal for the employer to assign all those attaining at least s to the high skilled job. Along EE the posterior probability that a worker is qualified is constant (except possibly at the end points).

The line WW(s) gives the fraction of workers that chooses to be qualified given any standard s. If the standard is 0, investing is useless since everyone will be promoted, and if the standard is 1, investing is again useless since no one will be promoted. Given our assumptions, the graph is single peaked.

In an equilibrium the firms and the workers are all optimizing. Thus, in an

---

[4]This figure corresponds to figure 2 in Coate and Loury(1993).

equilibrium the standard(s) that the firms set cause the workers to become qualified at precisely the rate which makes the standard(s) optimal. Clearly, this occurs at an intersection of the EE and WW curves. Thus, one equilibrium involves the firms setting a uniform standard of s_w in figure 1.

Consider an initial out of equilibrium standard. Coate and Loury posit the adjustment process in which workers respond optimally to this signal, firms then set the optimal standard given the proportion of workers who are qualified, workers respond to the new signal and so forth... This yields $\pi^{t+1} = G(\beta(s^*(\pi^t)))$, and the stability of an equilibrium is then defined with respect to this process.

One nice thing about their model, however, is that it is not necessary that firms have such a sophisticated understanding. Suppose that firms know nothing of Bayesian probability theory. A natural way for firms to behave is to raise the standard they have set if it turns out that marginal workers are qualified at too low a rate, and to lower the signal if marginal workers are qualified at too great a rate. If firms make gradual adjustments in this manner and workers respond optimally to these adjustments, all stable equilibria under Coate and Loury's process are stable under this process (made appropriately precise) but other equilibria may also be stable. More precisely, all equilibria to the right of $\bar{s}$ are either stable or unstable under both adjustment processes (for instance, s_u is unstable with respect to both processes). The equilibrium to the left of $\bar{s}$, however, may be unstable under the Coate and Loury adjustment process, but stable under the second process. I refer to this second process as the "natural adjustment" process and it is the process that I favor, though this has no qualitative effect on the results.

Now while a single standard of s_w is an equilibrium, there is also the possibility of a stable "discriminatory" equilibrium. The firms may set a standard of s_w for W's and s_b for B's (see figure 1). Given these standards W's, qualify at the rate π_w and B's qualify at the rate π_b. Firms may then justify the setting of a higher standard for B's by rightfully asserting that these higher standards are needed because B's tend to be less qualified. Of course, however, the only reason that B's tend to be less qualified is that a higher standard has been set for them.

It may seem that this is a situation ripe for affirmative action. However, Coate and Loury point out that an affirmative action program may have unintended consequences. Take a program which mandates that employers promote workers from both groups at the same rate. Although one way for employers to do this would be to set a standard of s_w for all workers, there may also be another way.

Consider the number or workers who are promoted at any standard. As the standard decreases in the range [s-,1], passing the test becomes easier and more workers choose to be qualified. Thus, the proportion of promoted workers passing the test is increasing. As the standard falls in the range $[0,\bar{s}]$, however, the number of qualified workers also falls. The net effect is ambiguous. In particular, it is possible that as s falls below s_w, the proportion of promoted workers initially falls, only to rise again later, so that at some standard $s_p < s_w$ (see figure 1), the rate at which workers are promoted is the same as at s_w[5].

Suppose that employers are forced to rapidly equate the rates at which B's and W's are promoted. Assume that the proportion of W's in the population is much greater than the proportion of B's. Firms will then hold the promotion rule for W's more or less fixed and adjust their promotion of B's. Given their initial negative beliefs, firms may well decide that the only way to promote a significant number of B's is by setting some very low standard close to 0. At a very lower standard, more than enough (unqualified) B's will be promoted. When firms then increase the standard for B's, they end up with a promotion rate equating standard of (approximately) s_p for B's. Consider also starting at a standard slightly below s_w for B's. At this standard, B's are not being sufficiently promoted. The natural adjustment on the part of firms is to lower the standard for B's, again leading to the standard s_p. Thus, employers may react to the affirmative action program by setting a low "patronizing" standard for B's which only serves to perpetuate, and may even aggravate, the skill disparity.

3. Discussion of the Coate and Loury Model

Although formally this model may be interpreted as one of hiring or promotions, the latter interpretation (which is the one used by Coate and Loury) seems better to me. In fact, the discriminatory equilibrium corresponds quite nicely to a story of "glass ceilings". Despite reaching the same standards as their co-workers, B's find that they are not promoted. This is exactly the complaint of those who assert they are victims of glass ceilings. Thus, many female professionals maintain that despite records comparable to those of their male counterparts, they are not promoted to the highest levels. Their bosses reply that the reason is that women are not dedicated enough to their job; for instance, they devote too much energy to their families. As the model shows, even if that is true, this behavior may simply be a rational response to the correct expectation that a female will probably not be promoted in any case[6].

Coate and Loury purport to show that despite this, an affirmative action program may not be called for. However, a closer look reveals that their analysis is less an argument against affirmative action, than an argument for a different type of affirmative action program than the one implicitly envisaged by them (as I read it). Instead of requiring instant and continuous equal promotion rates, suppose the affirmative action program requires that firms gradually increase their promotion of B's until the two groups are promoted in approximately equal proportions. Refer back to figure 1. What will happen if firms are forced to gradually lower their standards for B's? As the firms lower their standards the proportion of B's who

---

[5]Again, the reader should consult Coate and Loury(1993) for a fuller description of the conditions under which this possible.

[6]In this context, consider the following quote from Capital Business District Review, Aug. 24 1992, "But as long as more men than women, for whatever reasons, pursue their careers single-mindedly, say observers like Pendergast and McNamee Lochner's Buchanan, more men will make it to the corner office"(emphasis added).

become qualified increases. Initially, however, this increase is not enough to offset the increased ease of passing the test. The WW line lies below the EE line and there is grumbling on the part of firms as they lose money on the marginal workers. This grumbling continues until the firms are setting a standard of s_u. But notice what happens beyond this point. Now, the EE line lies above the WW line. The marginal B's have too high a probability of being qualified! The affirmative action program is no longer binding. The firms on their own will continue to lower the standard for B's until the uniform standard of s_w is reached. If for some reason, the firms overshoot, and set a standard slightly below s_w for the B's they will not misreact by lowering the standard further for B's. They are only required to hire the two groups in approximately equal proportions and so they will feel free to raise the standard for B's. The color-blind equilibrium of s_w will be reached. Thus, the right kind of affirmative action is beneficial, and the affirmative action program itself will only be temporary.

4. Disadvantaged B's

There is something odd about the affirmative action model of section 1. The question posed by Coate and Loury is essentially "Given that B's face a <u>higher</u> standard than W's, should employers be forced to lower this standard?" But this is not the question that has provoked such heated debate. Indeed, Coate and Loury's argument notwithstanding, and even leaving aside my reply, many opponents of affirmative action would no doubt answer yes to this query. But these opponents have another situation in mind. Namely, one in which members of the designated class fail to be hired or promoted in significant numbers when the standard they face is the <u>same</u> as the standard face by the majority. The question is then "Should B's be presented with a lower standard than W's in order to ensure that both groups are hired or promoted at similar rates?".

Thus in 1985, following a suit, the New York City Police Department promoted to sergeant African-American and Hispanic officers whose scores on the sergeant exam were lower than failing scores of their white counterparts. Similarly, some university affirmative action programs have led to the acceptance of minority members with lower SAT and other standardized test scores than those of non-minorities[7]. In both cases the affirmative action program entails a move away from an initial situation in which all groups face the same standard. Notice that in the model of section 1 such a move would not be necessary. If employers start out with a uniform standard, workers in both groups are promoted at the same rate[8].

---

[7] Although a university may not quite be a profit-maximizing firm, it can be viewed as receiving a benefit from admitting a qualified student and a loss from admitting an unqualified student. Thus this framework is an appropriate one. Also, the interpretation of the model as one of hiring and not hiring is appropriate in this section, and is the one that will be emphasized. The test may be of the amorphous type envisaged by Coate and Loury or a more objective measure.

[8] Coate and Loury (1993) does consider a variant without this property. I discuss this in section 4

Previously it was assumed that the two groups of workers, B and W were identical. Supporters of affirmative action, however, usually begin from the premise that all groups are not identical. In particular, for a variety of reasons members of minority groups are disadvantaged. I now incorporate this assumption into the model. In particular, I assume that a qualified B is likely to score lower on the test than a qualified W and that an unqualified B is likely to score lower than an unqualified W. For instance, given two equally hard working and intelligent students, it is natural to expect that the student from a poorer background attending a poorer high school, will perform less well on the SAT test. To the extent that a college is concerned with the industriousness and intelligence of the student, and not in the ability to perform well on the SAT test, the test score will be misleading. Similarly, an equally qualified worker from a disadvantaged background may do less well on an employment "test", as such a test inevitably will also test factors unrelated to job ability. For instance, an employer may judge the seriousness and reliability of an applicant based on an initial interview. One way for an applicant to come across as serious and reliable is to have spent a portion of his leisure time, say, at the library rather than partying. The interviewer's judgment of the applicant, however, will be imperfect at best and may well be racially biased, even if only inadvertently. Thus, a minority applicant may expect to "score" lower than a white counterpart who has spent an identical amount of time in the library[9]. Note, however, that nothing I have said prevents the employer from correcting for the interviewer's bias.

Let us now posit two sets of test density functions f_u and f_q for the W workers, and h_u and h_q for the B workers, with associated distribution functions F_u, F_q, H_u, and H_q. Since B's are disadvantaged, a B worker is less likely to score well than an equally qualified w. That is, for all θ $H_i(\theta) \geq F_i(\theta)$, i = u, q. For simplicity I assume that B's simply score proportionately lower than W's. Specifically, there exists an $0 < \alpha < 1$ such that $h_i(\alpha\theta) = 1/\alpha$ $[f_i(\theta)]$, and hence $H_i(\alpha\theta) = F_i(\theta)$. Mathematically, this has the effect of simply rescaling the equations and diagrams from the previous section[10]. α provides a measure of the degree to which B's are disadvantaged relative to W's, with lower α's indicating a greater disadvantage.

There are now two sets of EE and WW curves, one for each group, with $EE_b(\alpha s) = EE_w(s)$ and $WW_b(\alpha s) = WW_w(s)$. As in section 1 an EE_i curve may intersect a WW_i curve, i = w, b, in several places. However, this is not the focus of our interest here. For now, let us concentrate on the leftmost intersections. Refer to figure 2. First consider the "color-neutral" equilibrium where W's face a standard s_{cn} and B's face the lower standard αs_{cn}. Given these standards, workers in the two groups become qualified at the same rate, and pass the test at the same rate. The equilibrium is color-neutral in that the likelihood that a worker is qualified and promoted is not a function of the worker's color.

---

[9] Consider also the following assertion with respect to the aformentioned seargent's exam " 'If the test had been fair, Hispanics would have passed at the same rate as whites,' said Luis Salgado, president of the 2,000 member Hispanic Soxiety of the city's police department." United Press International, Sept. 14 1984.

[10] Notice that with this assumption, the test scores of B's is distributed over $[0,\alpha]$. A more general assumption would have $H_i(\alpha(\theta)\theta) = F_i(\theta)$, where $\alpha(\theta) < 1$ for $0 < \theta < 1$, and $\alpha(1) \leq 1$. All the results in the paper can be obtained with such a model in which $\alpha(1) = 1$, so that the B's test scores are also distributed over the entire interval [0, 1].

Suppose, however, that firms are not color-neutral, but rather are "color-blind". That is, they set the same standard for both groups. For concreteness, suppose that W's are a much larger fraction of the population than B's. First consider the situation in which firms are setting the uniform standard s_{cn}. W's and B's respond by becoming qualified at the rates π_{cn} and π_l, respectively. Then a W worker obtaining s_{cn} has the "right" probability of being qualified, but a B worker has too low a probability since the curve EE_b is above the WW_b curve at this point. Thus, the overall posterior probability that a worker is qualified is too low. The color-blind equilibrium will be a single standard slightly to the right of s_{cn}, say at s_{cb} in figure 2.

Is this truly an equilibrium? Yes, it may well be. Of course, as I will argue more fully later, a firm could increase its profits by adopting two different standards for the two different groups, but the firm need not realize this. That is, we as modelers, have divided the population into B's and W's, but the firm might not have done this. Put differently, one could also consider married workers vs non-married, those that commute vs those that do not, etc... It may be too costly for firms to try all these possibilities and so they may prefer to just set a single standard. Alternatively, there may be administrative or "socio-political" reasons for adopting a single standard.

Now while the standard is color-blind, the result is not color-neutral. The B workers are being hired or promoted at a lower rate than the W's. This occurs for two reasons. First of all, the "same" standard is harder to reach for the B's than the W's. Second of all, a smaller percentage of the B's choose to become qualified. As the next proposition indicates, if the B's are sufficiently disadvantaged, then in all color-blind equilibria B's will be less qualified and will be hired at a lower rate than W's.

Formally, a color-blind equilibrium is a single standard s such that when workers optimize with respect to the standard, it is optimal for firms to promote all workers who have achieved this standard, given the constraint that promotion is independent of color. Let λ be the fraction of W's in the population. We have the following proposition:

Proposition 1: (color-blind is not color-neutral) There exists a $\bar{\lambda}$ and $\bar{\alpha}$ such that for all $\lambda > \bar{\lambda}$ and for all $\alpha < \bar{\alpha}$ in any color-blind equilibrium B's are qualified at a lower rate than W's and pass the test at a lower rate.

Proof: Let s_1 be the leftmost intersection of EE_w with WW_w. Now for $0 < \alpha < 1$ the curves WW_w and WW_b intersect at $s=0$ and one more point. As α varies from 0 to 1 this second intersection increases from 0 to $\bar{s} = \text{argmax}\{WW_w(s)\}$. Choose $\bar{\alpha}$ so that this second intersection is to the left of s_1 (refer to figure 2, where $s_1 = s_{cn}$). Then at s_1, the WW_b curve lies below the WW_w curve. For λ close to 1, the color-blind equilibria occur close to the intersections of the EE_w and WW_w curves. At all these intersections the WW_b curve lies below the WW_w curve, so that B's are less qualified than W's and pass the test at a lower rate.

Q.E.D.

Note that if s_1 is to the left of $\bar{s}$, then for α close enough to 1 B's will be more qualified than W's in the leftmost equilibrium (but not in the other equilibria). However, if α

is very close to 1 the difference in qualification rates is not great. A more interesting comparison along these lines is given by the next proposition.

We have been considering employers with two jobs, a skilled job and an unskilled job (which may also be interpreted as not hiring the worker). Skilled jobs, however, vary in the degree of skill required. In a low-skilled job an unqualified worker may be almost competent so that the loss, x_u, to misassigning such a worker is small, whereas in a high skilled job it is more important that only qualified workers be promoted (or hired). Therefore, we can think of $r = x_q/x_u$ as measuring the skill level of a job, with smaller values of r indicating higher skilled jobs. The next proposition indicates that color-blind equilibria work against B's in high skilled jobs. More precisely, B's will be less qualified than W's for high skilled jobs, and will be hired at a lower rate in all interior color-blind equilibria. For low skilled jobs, B's will be more qualified than W's in the leftmost equilibrium -- they may or may not be hired at a higher rate.

<u>Proposition 2</u>: (skill discrimination) There exists a $\bar{\lambda}$ and r* such that for all $\lambda > \bar{\lambda}$ and $r < r^*$ in any interior color-blind equilibrium B's are less qualified than W's and are hired at a lower rate. For $r > r^*$ B's are more qualified at the leftmost color-blind equilibrium -- they may or may not be hired at a higher rate.

<u>Proof</u>: Let (π', s') be pairs along the WW_w curve. For $0 < s < 1$, points along the EE_w curve satisfy $r = [(1-\pi)/\pi]f_u(s)/f_q(s)$ so that at intersections of the two curves $r = [(1-\pi')/\pi']f_u(s')/f_q(s')$. As r decreases from infinity, the first intersection of EE_w and WW_w increases smoothly from $s = 0$ to $s = \bar{s}$; as r decreases further, this intersection continues to increase to 1, although possibly discontinuously. Let $s_1(r)$ be this first intersection. We now choose r* so that the left-most color-blind equilibrium is at the standard s* where WW_w and WW_b intersect. In fact, r* is such that $s_1(r)$ is a little to the right of s*. For then at $s_1(r)$, the marginal W's are properly qualified, but the marginal B's are qualified at too great a rate (this follows from the fact that at $s_1(r)$ the WW_b curve lies above the EE_b curve (which lies below the EE_w curve)). Hence, the leftmost equilibrium is a little to the left of $s_1(r)$, and for appropriate choice of r*, it will be exactly at s*. Note that at s*, B's and W's are qualified in the same proportion, but W's are hired at a higher rate. Now as r decreases, the leftmost color-blind equilibrium moves to the right. Thus, for $r < r^*$, in all color-blind equilibria B's are less qualified and pass at a lower rate than W's. For $r > r^*$, B's are more qualified than W's in the leftmost color-blind equilibrium. They may or not pass at a higher rate[11]. (Note that if B's are passing at a high enough rate, the color-blind equilibrium may be closer to the intersection of the EE_b and WW_b curves, than the intersection of the EE_w and WW_w curves.)
Q.E.D.

---

[11] If the proportion of workers who pass the test is a mononotonically decreasing funtion of the test score, than B's will pass at a lower rate than W's, even though they are more qualified. If this function is not monotonic (which is also the necessary condition for the patronizing equilibrium of section 2), then B's may pass at a higher rate. Additional assumptions are required to guarantee the existence of stable equilibria of the various types described.

The comparative static in proposition 2 has been "partial equilibrium" in nature. That is, we have been looking at the effect of different skill levels, holding all other parameters fixed. It is likely, however, that a high skill job pays a high wage as well. At the same time, the investment cost necessary for becoming qualified for a high skill job may also be high. The net effect of all these differences is not obvious. An examination of the proof of proposition 2 reveals that the key to determining whether B's are more or less qualified than W's in equilibrium is the level of the standard. When the standard is low, B's are more qualified than W's; when the standard is high B's are less qualified than W's. To the extent that a high skill job is characterized by a relatively high standard and low pass rate, B's will be less qualified and pass at a lower rate than W's in these jobs.

Thus, the color-blind equilibria may work against B's. But is there a need for an affirmative action program? Won't some firms realize that the test is biased and lower their standard for B's in order to pick up low scoring but qualified B's? Even if firms do not realize this on their own, should not the mere suggestion of an affirmative action program be sufficient to nudge them on their way to the color-neutral equilibrium? The answer is no, not if B's start out sufficiently disadvantaged[12].

Refer again to figure 2. Starting from the color-blind equilibrium s_{cb} suppose the firms become color-conscious. What will happen? Notice that in the equilibrium the marginal B workers are failing at too great a rate (the EE curve is above the WW curve). The natural reaction on the firms' part will be to increase the threshold for B's. In fact, starting at the color-blind equilibrium any firm can increase its profit by increasing the threshold for B's and decreasing the threshold for W's. When all firms do this, an adjustment on the part of the workers ensues and a further adjustment on the part of the firms and so forth until the stable equilibrium pair (s_{cn}, s_{cc}) is reached. In this equilibrium the disadvantaged B's face a higher standard than the W's, rather than the lower color-neutral standard. Of course, B's are hired at a lower rate than W's.

Return again to the color-blind equilibrium. Suppose that a sophisticated firm realizes that the test is biased against B's and decides to lower its standard rather than raise it. What will happen? Since the firm is only one of many, this change will have little effect on the investment behavior of the B workers. Therefore, the reduction will cause the firm to lose money. Even if all other firms lower their standards for B's by a little, the individual firm will do better by raising its standard. Lowering standards towards the color-neutral level is rational only if a sufficient number of firms move together, and lower their standards by a sufficiently large amount. Something like a government mandated affirmative action program is needed in order to move the firms to the Pareto improving color-neutral equilibrium. In contrast, raising the standard for B's, is rational for the individual firms moving individually or together.

The situation as depicted in figure 2 is only one of many possibilities. Consider

---

[12] Spence(1974) also observes that in an unequal world, color-blind signalling may be unfair. In Spence's model, however, individual firms can costlessly lower signalling requirements for disavantaged groups.

figure 3. Now in the color-blind equilibrium s_{cb}, the marginal B workers are succeeding at too great a rate. If firms become color-conscious they will lower the standard for B's, increase the standard for W's, and continuing in this way move to the color-neutral equilibrium on their own. Here there is no need for an affirmative action program. Naturally enough, the need for an affirmative action program depends upon empirical parameters.

As the next proposition indicates when B's start out sufficiently disadvantaged, this disadvantage will persist. Starting from the color-blind equilibrium, firms will go to a color-conscious equilibrium with a standard for B's which may be higher or lower than the standard for W's but, in either case, will not be a color-neutral equilibrium. B's will remain underemployed.

Recall that in the natural color-conscious adjustment process firms lower standards for a group whose marginal members are found to be qualified with too great a probability and raise standards for a group whose marginal members are found to be qualified with too low a probability. A color-conscious equilibrium is a pair of standards s_b and s_w such that $EE_b(s_b) = WW_b(s_b)$ and $EE_w(s_w) = WW_w(s_w)$.

<u>Proposition 3</u>: (persistent disadvantage) There exists a $\bar{\lambda}$ and $\bar{\alpha}$ such that for all $\lambda > \bar{\lambda}$ and for all $\alpha < \bar{\alpha}$, starting from any stable color-blind equilibrium, the natural adjustment process ends in a non-neutral color-conscious equilibrium in which B's are less qualified and hired at a lower rate than W's. They may face a higher or lower standard than W's. If α is small enough then no B's are hired.

<u>Proof</u>: Consider the intersections of the EE_w and WW_w curves, starting with the first stable intersection. Let $s_1 < \ldots < s_n = 1$ be the standards corresponding to these intersections, with s_1 being the first stable intersection. s_2 is then an unstable intersection. Let $\bar{\alpha} = s_1/s_2$. Then the curves EE_b and WW_b intersect at s_1, as shown in figure 4. The standard s_1 forms part of a color-blind equilibrium. Now consider an α slightly less than $\bar{\alpha}$. Then at s_1, W's are properly qualified, but B's are underqualified. For λ sufficiently close to 1, the color-blind equilibrium near s_1 is at a standard slightly higher than s_1. In this equilibrium, W's are overqualified and B's are underqualified (see figure 2, with $s_{cn} = s_1$, $s_u = s_2$, s_{cb} = the color-blind equilibrium). If employers now become color-conscious they will lower the standard for W's and increase the standard for B's. The natural adjustment process will lead to a color-conscious stable equilibrium with a standard of s_1 for W's and a standard of $\alpha s_3 > s_1$ for B's. Note that in the equilibrium B's are less qualified and pass the test at a lower rate than the W's.

Now refer to figure 5. Here α is again less than s_1/s_2, but this time α is slightly less than s_1/s_3 (with $s_3 < 1$). In the color-blind equilibrium close to s_1, marginal B's are overqualified. The color-conscious adjustment now involves a lowering of standards for the B's, but only to $\alpha s_3 > \alpha s_1$. Again B's are less qualified than W's and are promoted at a lower rate.

The analysis so far has been done with respect to the color-blind equilibrium near s_1. If $\bar{\alpha} = \min \{s_i/s_{i+1} : s_i \text{ is stable}\}$, then the same analysis establishes the proposition starting from any stable color-blind equilibrium. Finally note that if $\alpha < s_1$ no B's are hired in the color-blind equilibrium, and this equilibrium is also a color-conscious equilibrium (the case

$\alpha < s_1$ also quickly establishes much of the proposition. However, this case also seems to suggest that there is only a problem when no B's are being hired; the rest of the proof shows this is not the case).

Q.E.D.

On the other hand, if α is close enough to 1 the color-conscious adjustment process will lead to a color-neutral equilibrium. There will be no need for an affirmative action program, but then there will not appear to be such a need either. When B's are consistently underhired or underemployed some kind of affirmative action program will be necessary.

5. What kind of disadvantage?

Although Coate and Loury's main model does not account for B's being hired at a lower rate than W's when the two groups face identical standards, they do consider a variant with this property. They suppose that due to disadvantages, B's have a different cost distribution than W's. Specifically, they posit two cost distribution functions with $G_b(c) < G_w(c)$ for $0 < c \leq ß(s\text{-})$. This alternative, and not implausible, way of describing the disadvantaged position of minority members has different implications than the shift in test distribution function posited in the previous section. Graphically, the shift in G(c) causes the WW curve to move down for the B's, as in figure 6. At a color-blind equilibrium, s_{cb} or t_{cb}, or at a color-concious equilibrium, (s_w, s_b) or (t_w, t_b), B's are hired at a lower rate than are W's. Note that there is now no color-neutral equilibrium in the sense of section 3. Suppose that firms try to increase the proportion of hired B's somewhat by lowering the standard for them. If the initial color-blind equilibrium is to the right of $\bar{s}$, then this will induce an increase in the number of B's who get qualified[13]. However, if the initial color-blind equilibrium is to the left of $\bar{s}$, this reduction in the standard has an ambiguous effect; depending upon the parameters the qualification disparity may increase.

To evaluate the merits of affirmative action it is of some importance to determine the nature of the disadvantage suffered by minority groups. Both Coate and Loury's assumption and mine imply that when a uniform test standard is set, members of disadvantaged groups will pass at a lower rate than others. Coate and Loury's assumption implies further that those disadvantaged group members who pass will have a lower probability of being qualified than others. My assumption is less categorical on this point, although if the group is disadvantaged enough this will still be true. There is, however, an implication on which the two assumptions necessarily differ. Under Coate and Loury's assumption, the expected test score of an individual given his (ex post) degree of qualification is independent of his race. Under my assumption, this conditional expectation is lower for members of disadvantaged groups. This may provide a way of distinguishing empirically between the two assumptions. Doubtless, there is some truth to both of them. At the same time, there are still other ways in which minority groups may suffer, including pure discrimination unrelated to any real differences.

---

[13] The reader should not be misled by the diagrams. It is perfectly possible that all equilibria are to the right of $\bar{s}$.

6. Some Further Empirical Considerations

The proposition that some employment and admission tests are biased against minority members is, in principle, testable. There are certain difficulties, however. Although one might be tempted to evaluate the biasedness of tests by estimating the probabilities that agents in different groups succeed given their test scores, this is misguided. These probabilities also pick up the effect of the equilibrium number of agents who have chosen to become qualified. The proper parameter to estimate for different groups is the expected test score, given the qualification of an agent.

If two groups are qualified in the same proportion, then the assumption that a test is biased against one group implies that the test score will underpredict the performance of that group. However, when these proportions differ, this is no longer true. Examining figure 2, we see that even if equally qualified B's score less well on a test, the test score might overpredict the performance of B's (this happens, for example, when a common standard of s_{cb} is used). Thus, even if, for instance, the SAT is biased against a certain minority group, the relationship between the regression of grade point averages against SAT scores for the disadvantaged group and the majority gtroup is unclear. It could very well happen, for example, that the minority regression line lies above the majority regression line for low test scores but lies below the majority line for high test scores; these latter scores being mainly "statistical errors" from unqualified students.

Suppose that it was determined that at universities with affirmative action programs, minority students admitted with low SAT scores performed significantly less well than other students. What would this indicate? It is not clear. This discrepancy could be less a reason to oppose affirmative action, then an argument for a more extensive affirmative action program. After all, only a sufficiently broad program can induce a positive investment response on the part of prospective students.

Determining which agents are qualified is itself no trivial matter. The difficulties in doing a proper survey of workers is obvious. Consider also colleges and the SAT test. While there is some controversy over exactly what it is the SAT measures, at the very least it measures the ability to do well on a test. Success in college is largely measured by the ability to do well on tests. Thus, even if the SAT is in fact biased with respect to the measurement of "true ability", it may appear less so when evaluated as a predictor of perceived ability. Of course, a college may place a high value on success in class tests per se, as may students. Hence, the distinction between being truly qualified and appearing to be qualified may not be that great.

On the other hand, recall the promotion of low-scoring minorities to sergeant in New York city. This settlement was reached after city officials determined that they could not establish that the sergeant's exam was job related. This suggests that the exam functioned as a pure signal where the score on the test per se was of little significance. Thus if the test was indeed biased, then promoting low scoring minorities would have no cost associated with it (given an appropriate supply response)[14].

Proposition 2 suggests the following interesting finding. Under color-blind employment practices few minority members are promoted to high-skilled jobs, and the proportion of qualified workers in the general pool is low. In some low-skilled jobs, however, relatively more minority workers are promoted and the proportion of minority workers who are qualified for these jobs is greater than the proportion of qualified white workers, both in the general pool and among those who are promoted.

7. Conclusion

Affirmative action is at best a temporary solution. Surely it is preferable to raise the performance of disadvantaged group members on employment and admission tests, than to lower the standards for these groups. In the meantime, however, affirmative action may be a valuable tool. Some commentators have suggested that inner city youth suffer from a lack of "correct" values; apparently the fault lies not in society but in these youths' own unwillingness to study more assiduously. However, it is not clear why they should. A hard-working inner city student still has a much smaller chance of gaining admission to college or employment in a good job than does a more privileged student. Affirmative action may be a way of providing the right incentives to the disadvantaged. Color-blind behavior on the part of firms and schools is not color-neutral when society is neither.

[14] It should be noted that the city did not admit that their tests were biased. Nonetheless, a new "unbiased" test was designed with the help of a paid consultant. Minority members fared no better on this test. Presumably, this test was designed to be free of elements which appeared to disfavor minority members. However, to my knowledge no systematic study of the on-the-job performance of low scoring minority mebers that were promoted has been undertaken. Such a study would be needed to see to determine if the purportedly unbiased test was biased in the sense used in this paper.

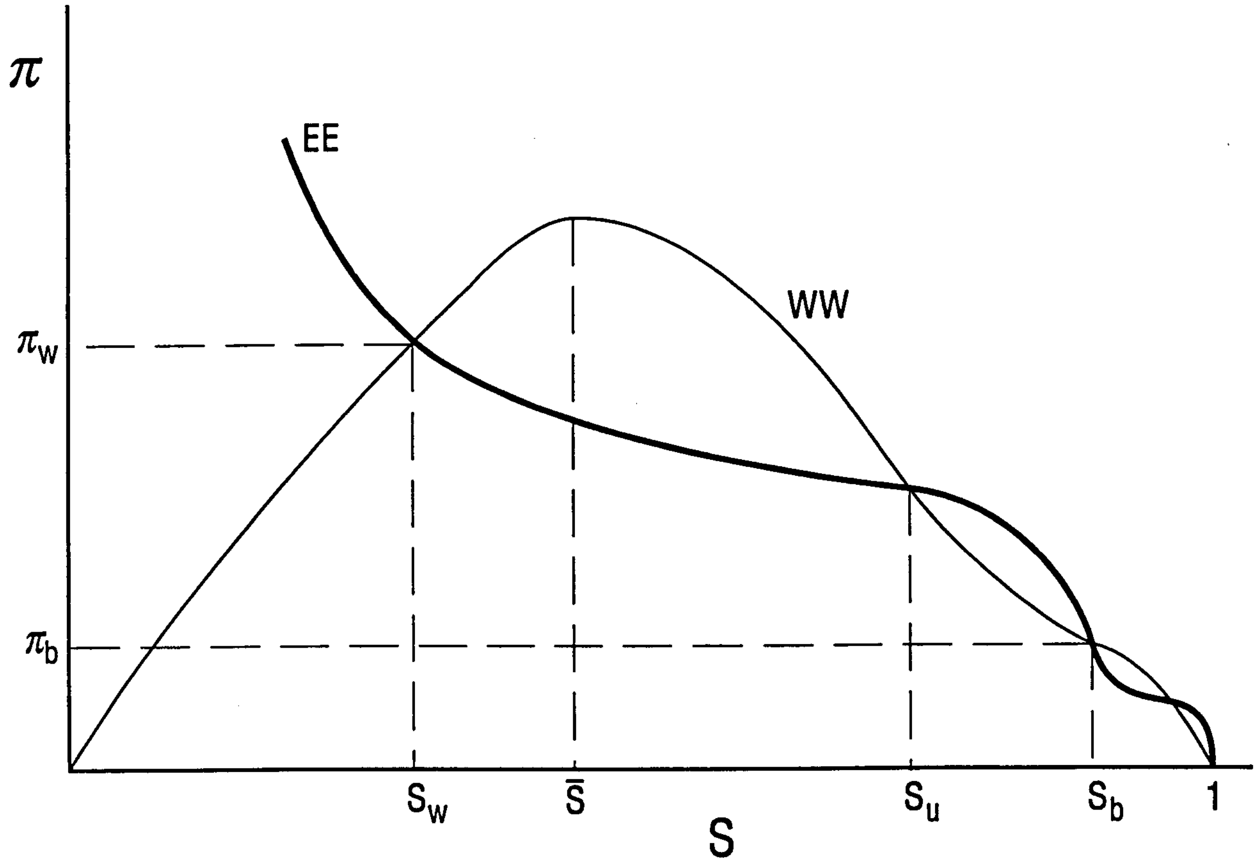

Figure 1

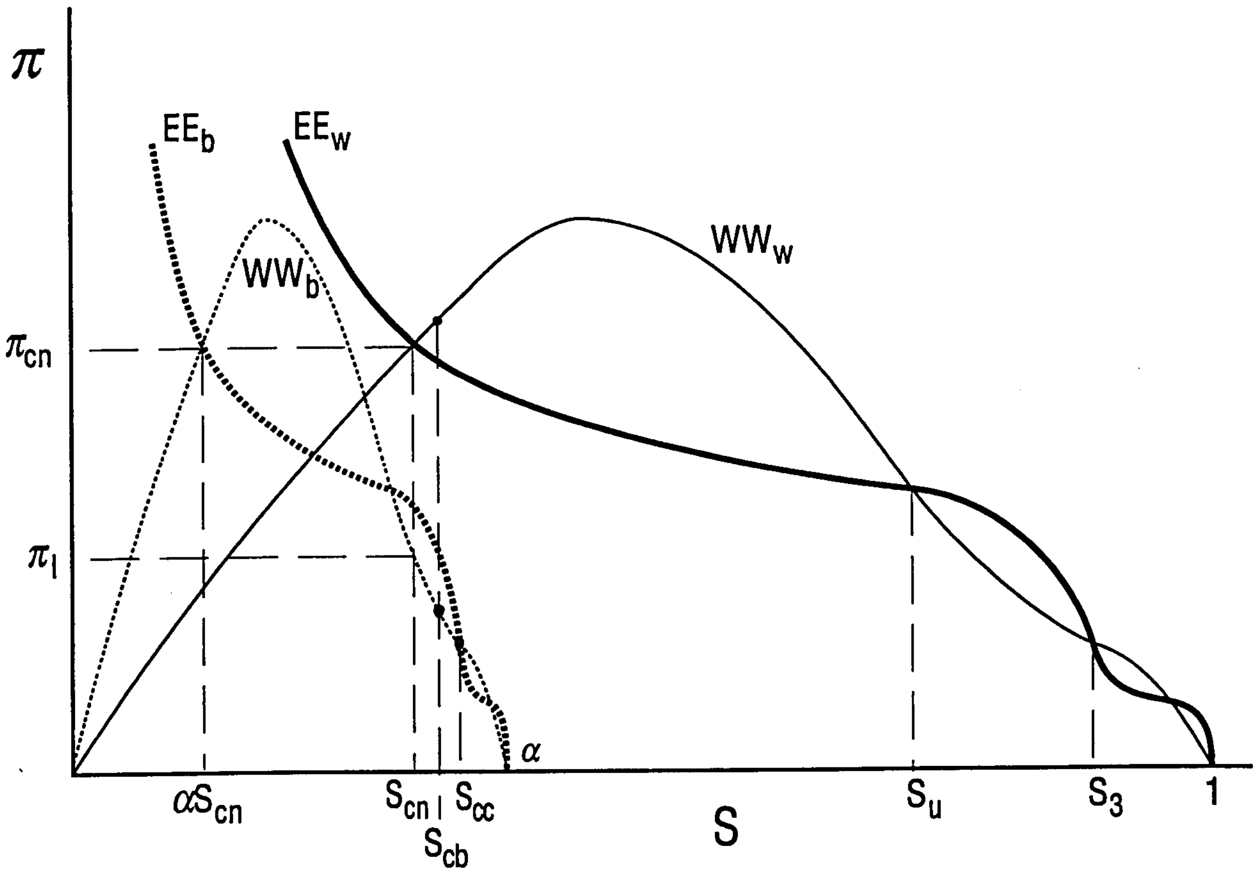

Figure 2

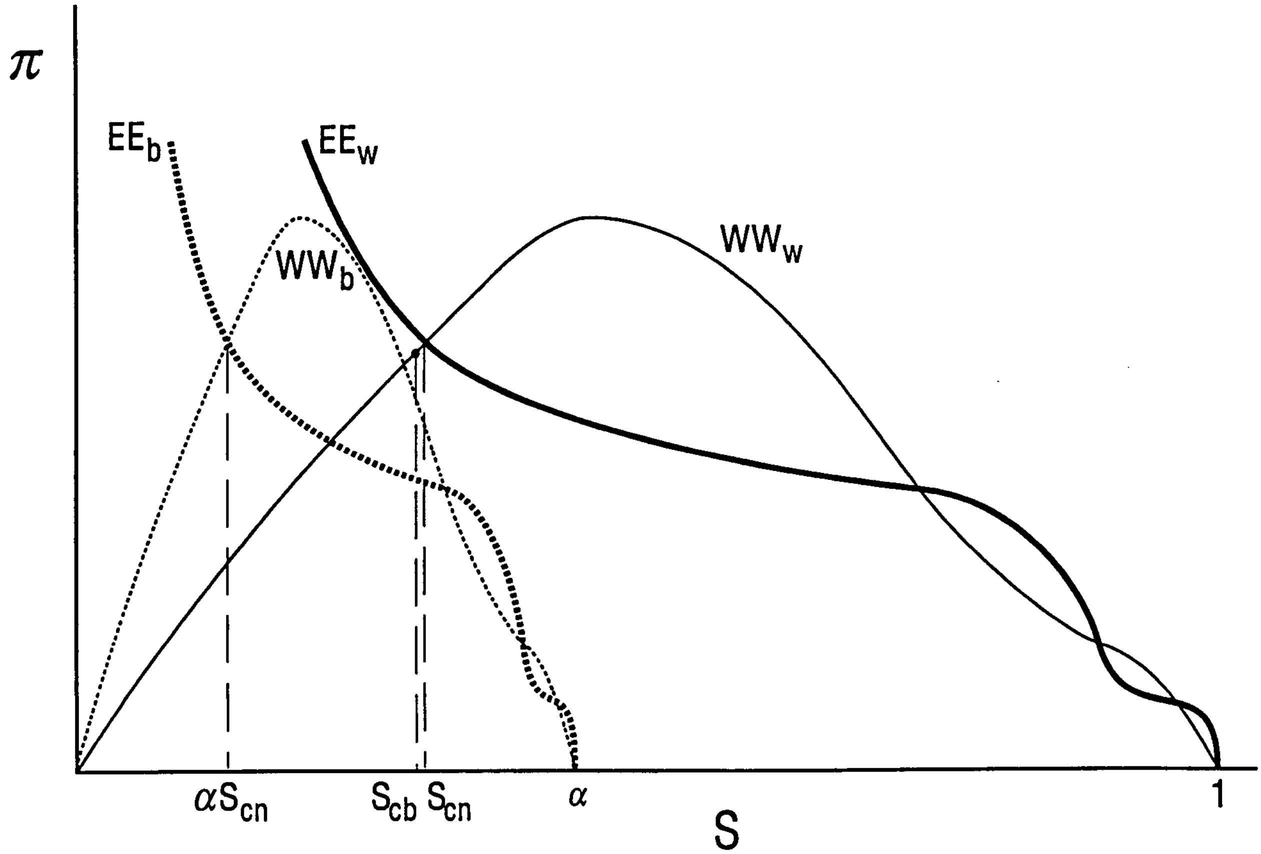

Figure 3

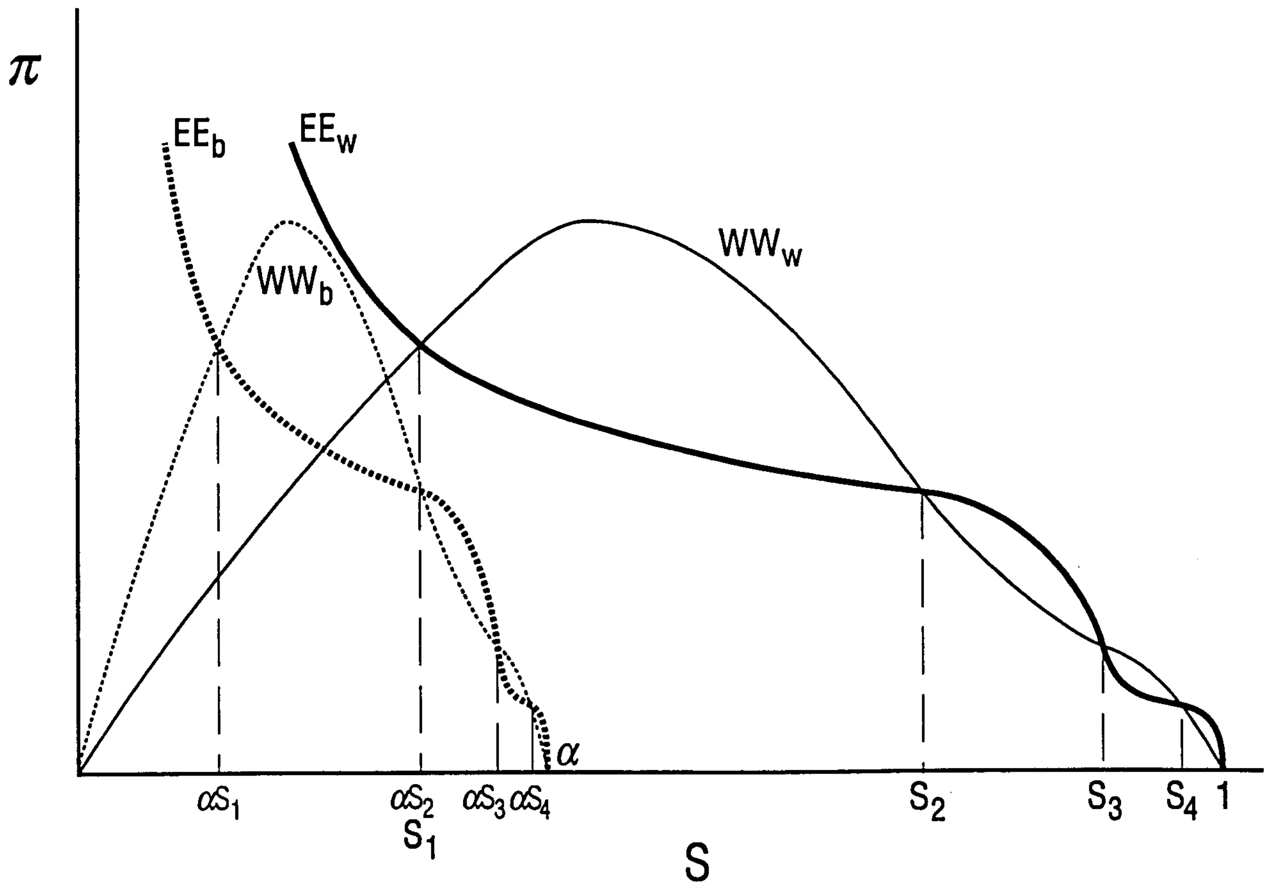

Figure 4

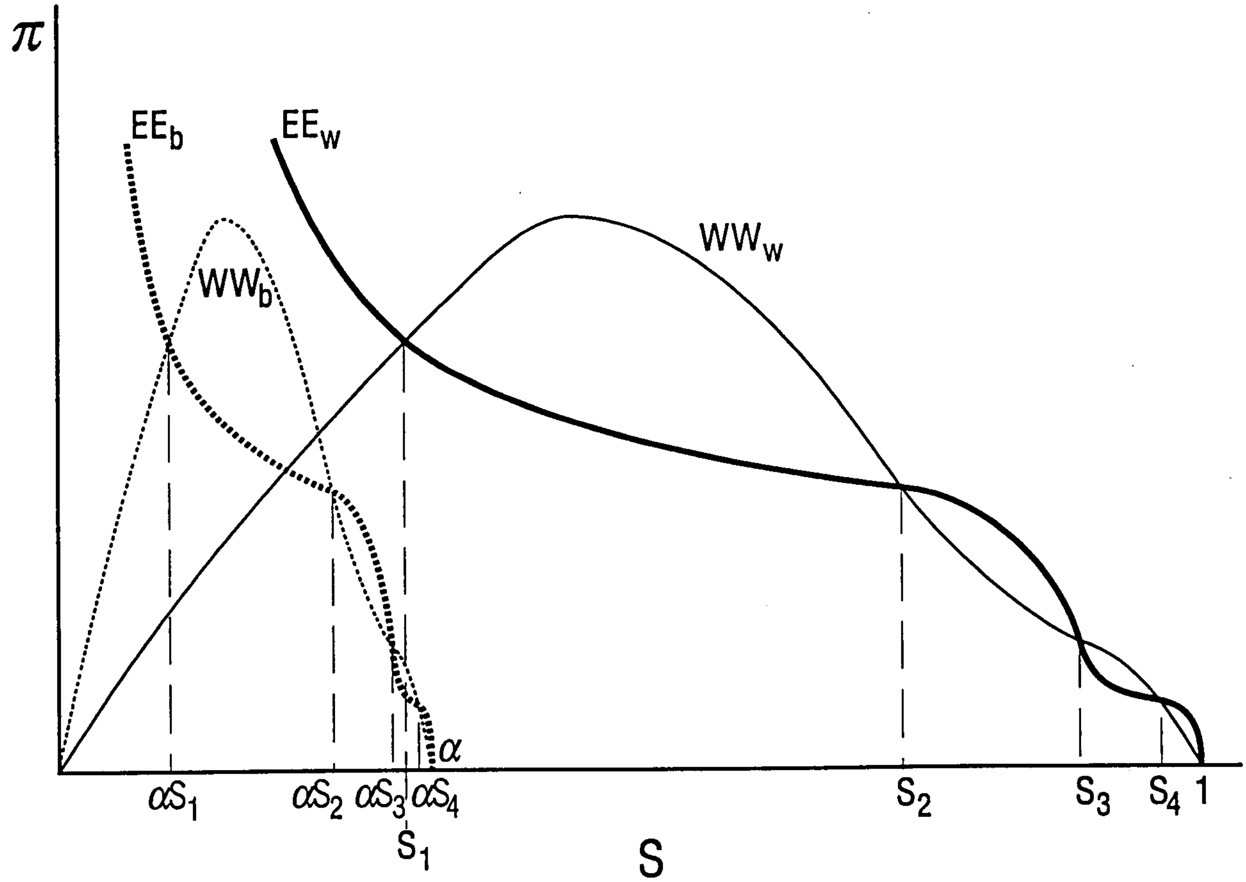

Figure 5

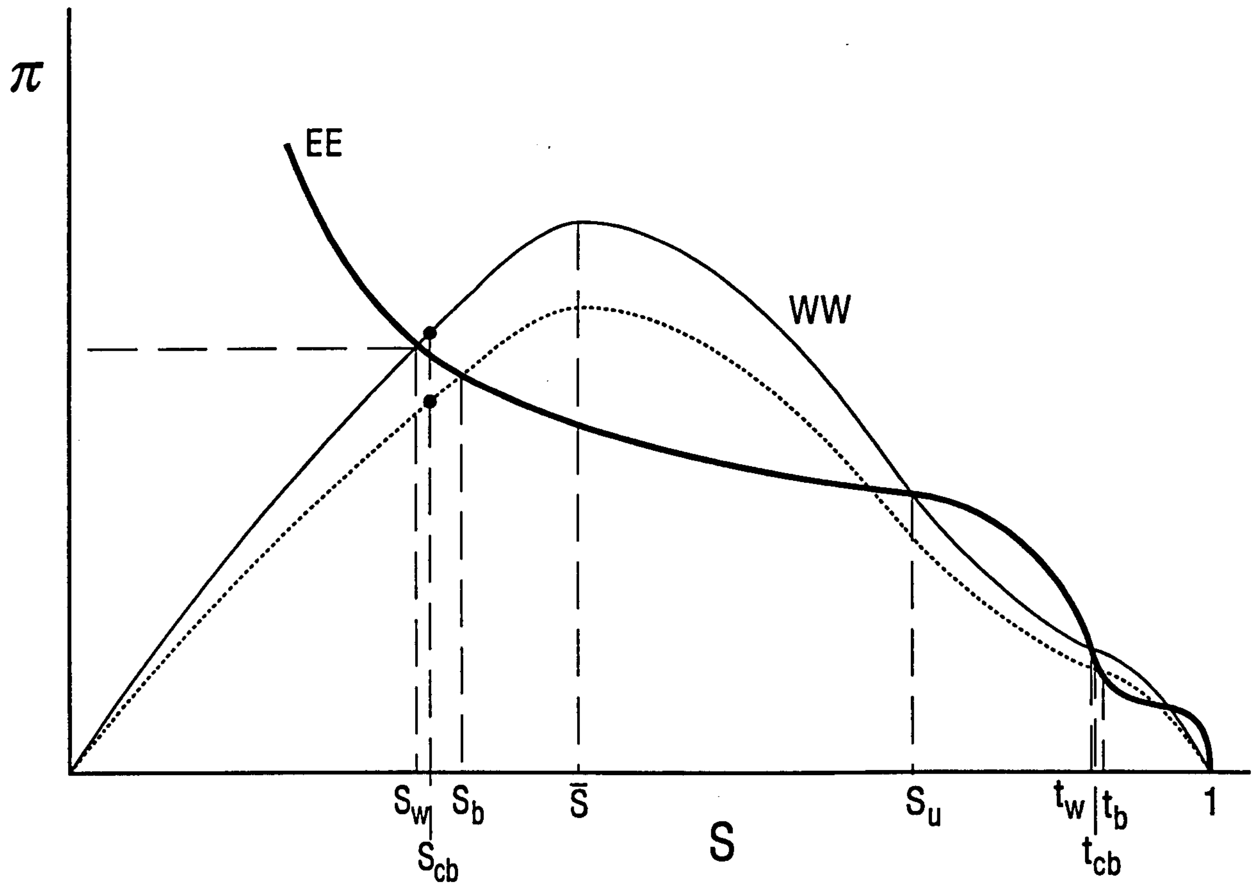

Figure 6

References

Arrow, Kenneth J., "The theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds. Discrimination in Labor Markets, Princeton, NJ: Princeton University Press, pp. 3-33

Coate, Stephen and Loury, Glenn C., "Will Affirmative-Action Policies Eliminate Negative Stereotypes?," American Economic Review, December 1993, 1220-1240.

Statistical Abstract of the United States 1993, United States Department of Commerce.

Spence, Michael A., Market signaling: Information Transfer in Hiring and Related Screening Processes, Cambridge, MA:Harvard University Press, 1974.

Welch, Finis, "Employment Quotas for Minorities," Journal of Political Economy, August 1976, 84, S105-139.