Columbia University

Department of Economics
Discussion Paper Series

Statistical Discrimination with Neighborhood Effects: Can Integration Eliminate Negative Stereotypes?

Shubham Chaudhuri
Rajeev Sethi

Discussion Paper No.: 0304-09

Department of Economics
Columbia University
New York, NY 10027

November 2003

Statistical Discrimination with Neighborhood Effects: Can Integration Eliminate Negative Stereotypes?*

Shubham Chaudhuri
Department of Economics
and SIPA
Columbia University
(sc301@columbia.edu)

Rajiv Sethi
Department of Economics
Barnard College
Columbia University
(rs328@columbia.edu)

November 2003

Abstract

We introduce neighborhood effects in the costs of human capital acquisition into a model of statistical discrimination in labor markets. This creates a link between the level of segregation and the likelihood and extent of statistical discrimination. As long as negative stereotypes persist in the face of increasing integration, skill levels rise in the disadvantaged group and fall in the advantaged group. If integration proceeds beyond some threshold, however, there can be a qualitative change in the set of equilibria, with negative stereotypes becoming unsustainable and skill levels in both groups changing significantly. This change can work in either direction: skill levels may rise in both groups, or fall in both groups. Which of these outcomes arises depends on the population share of the disadvantaged group, and on the curvature of the relationship between neighborhood quality and the costs of human capital accumulation.

Keywords: statistical discrimination, neighborhood effects, human capital spillovers

JEL codes: J71, J24, D82

*This material is based upon work supported by the National Science Foundation under Grant No. SES-0133483.

1 Introduction

Stereotypes — beliefs about some unobservable (or yet to be observed) trait of an individual based upon his membership in an identifiable group — are ubiquitous in economic and social interactions. Ever since the pioneering work of Arrow (1972, 1973) and Phelps (1972) it has been recognized that such stereotypes do not depend upon an active taste for discrimination and can be self-perpetuating even in the absence of any intrinsic group differences. In these early papers, and in the subsequent literature on what has come to be termed statistical discrimination, the stability of stereotypes arises because relevant traits are only imperfectly observable at the individual level. This allows unequal treatment based on *ex-ante irrelevant* group identities to emerge as an equilibrium outcome.

The contributions of Arrow and Phelps have spawned a considerable literature that explores whether the well-documented disadvantages in labor market outcomes experienced by individuals belonging to certain groups (such as racial and ethnic minorities or women) might be explained by the prevalence of negative stereotypes amongst employers. There are two strands in this literature. One follows Phelps (1972) in assuming that individual measures of productivity are noisier for members of disadvantaged groups.[1] The other, derived from Arrow (1973), does not assume such a difference. Instead, the key insight is that a stereotype can influence the behavior of those subject to it in ways that cause the belief to become self-fulfilling. So, for instance, if individuals from certain groups perceive that employers hold negative stereotypes about their capabilities, and are hence less likely to treat them favorably, their incentives to acquire human capital are diminished, thereby reinforcing and perpetuating the employer beliefs.[2]

The view of human capital acquisition that is implicit in these latter papers is one of a largely individual and autonomous process. The treatment that individuals anticipate in the labor market does influence their perceived benefits of acquiring human capital, but the costs of human capital acquisition are assumed to be exogenously given at the individual level. This assumption goes against the view, evident in both a range of theoretical papers as well as in a number of empirical analyses, that human capital investment decisions are subject to a variety of interpersonal externalities.[3] For instance, it is widely held that both the pecuniary and non-pecuniary costs of acquiring

[1] Aigner and Cain (1977), Borjas and Goldberg (1979), Lundberg and Startz (1983) all present models of statistical discrimination in labor market settings that share this feature.

[2] Coate and Loury (1993) and more recently, Moro and Norman (2003a, 2003b) are examples of papers in this vein.

[3] Consider, for example, the following statement in Lucas (1988): "Human capital accumulation is a *social* activity, involving groups of people in a way that has no counterpart in the accumulation of physical capital." Human capital-

human capital vary with the social and economic milieu in which an individual operates, and in particular, with the human capital of others with whom the individual associates, whether in the family, the neighborhood, at school, or in the workplace. And the usual presumption is that these interpersonal externalities take the form of positive complementarities — the higher the levels of human capital of others in a group, the lower the costs of human capital acquisition are for each individual within it.

If, indeed, local complementarities — what we loosely term neighborhood effects — are important, the costs of human capital acquisition at the individual level can clearly no longer be taken to be exogenous. Instead, these costs, and how they vary across groups, will be determined endogenously in equilibrium along with labor market outcomes and human capital levels. This has two important implications. First, the assumption, typical in the existing statistical discrimination literature, that the within-group distribution of human capital acquisition costs is the same for all groups can no longer be sustained, and the hypothesis that groups are *ex-ante* identical needs to be recast at a more primitive level.[4] Second, with the possibility of local human capital spillovers, the level of segregation in the economy becomes salient. That is because the extent to which individuals from different groups interact directly determines whether local human capital spillovers are confined within the boundaries of a group or extend across groups; and that in turn potentially influences whether group-level asymmetries in human capital investment levels and labor market outcomes can be sustained in equilibrium. Neither of these implications has been addressed in the existing literature on statistical discrimination. In this paper we take a step towards filling that gap.

based externalities, of course, feature prominently in endogenous growth models. Peer group effects are also considered in several models of endogenous community formation (see, among others, Arnott and Rowse (1987), De Bartolome (1990) and especially Benabou (1993, 1996a, 1996b)). On the empirical front, numerous studies provide evidence suggestive of a variety of peer group and neighborhood effects, among them, Coleman et al. (1966), Summers and Wolfe (1977), Henderson, Mieszkowski and Sauvageau (1978), Borjas (1995), and Ioannides (2001). However, there remain important concerns about the robustness and significance of this evidence. Manski (1991) expresses doubts that such endogenous social effects can plausibly be identified using the sorts of data that have typically been available, and Kremer (1997) questions the empirical importance of such effects in explaining stratification.

[4] The hypothesis that group differences are not innate but result instead from disparate treatement in equilibrium has been termed the axiom of *anti-essentialism* by Glenn Loury (2002). This is reflected in the *ex ante* equality of cost functions across groups in standard statistical discrimination models. While we allow cost functions to differ across groups in equilibrium as a result of endogenous differences in neighborhood quality, the mapping from neighborhood quality to cost functions is held to be the same for each group. Thus the anti-essentialist hypothesis is maintained throughout this paper.

We do so by incorporating neighborhood effects into an otherwise standard model of statistical discrimination. This allows us to explore the links between the level of segregation in the economy (which we take to be a primitive) and the likelihood and extent of statistical discrimination. In the presence of neighborhood effects, there are potential human capital spillovers from the advantaged group to the disadvantaged. The extent to which such spillovers arise depends on the degree of intergroup contact at the neighborhood level. Increasing integration therefore tends to lower the costs of human capital acquisition in the disadvantaged group while raising those in the advantaged group. If integration proceeds far enough, however, it may result in a qualitative change in the set of attainable equilibria, making it impossible for negative stereotypes to be sustained. In this case a shift to symmetric treatment of groups is triggered, resulting in rapid and significant changes in skill levels. While this tends to equalize labor market outcomes across groups, it can do so in one of two quite distinct ways. Depending on the population share of the disadvantaged group and the strength and curvature of neighborhood effects, the resulting shifts can leave both groups better off or both groups worse off than under the status quo. This suggests that under certain circumstances, both the advantaged and the disadvantaged group may be supportive of vigorously pursued integrationist policies, while under other circumstances, both may be opposed. More generally, the pursuit of integration, whether in the schools or in housing markets, may have significant and identifiable equity and efficiency effects that extend well beyond the specific arenas in which integrationist policies are attempted.

Our work is related to and complements a number of papers in several related areas. First and most obviously, the paper is related to the important contributions of Coate and Loury (1993) and Moro and Norman (2003a). These papers examine whether concerted public action in the form of temporary affirmative action programs can eliminate self-confirming discriminatory stereotypes. We investigate the related and as yet unexplored question of whether, starting from a situation where statistical discrimination is prevalent, public efforts towards greater integration of previously segregated groups can eliminate negative stereotypes in labor markets.

A related literature addresses the efficiency implications of statistical discrimination. Schwab (1986) is an early example while Norman (2003) provides the most recent and thorough analysis. The central trade-off emphasized in Norman (2003) is the trade-off between the cost inefficiencies implied by statistical discrimination — individuals in the disadvantaged group who have low costs of human capital acquisition choose not to invest, whereas higher cost individuals in the advantaged group do — and the gains from better matching (of qualified individuals to jobs that require high levels of human capital and unqualified individuals to jobs that do not) possible under statistical discrimination. We introduce an additional channel through which the elimination of negative stereotypes may affect efficiency. And that is through local human capital spillovers and whether these spillovers operate asymmetrically on individuals with high and low levels of human capital.

We show that the likelihood, extent and efficiency implications of statistical discrimination depend critically on the strength and direction of any asymmetries in the impact of neighborhood effects. These results echo those in the literature on endogenous community structure and stratification — the work of Benabou (1993, 1996a, 1996b) being a prominent example — where such asymmetric effects are shown to matter for the extent and efficiency implications of stratification. The paper also complements the work of Sethi and Somanathan (2002), which explores the consequences of income inequality for residential segregation. We take the converse approach, starting with a given level of segregation, and explore the implications for inequality in the form of statistical discrimination.

Finally, Moro and Norman (2003b) criticize standard statistical discrimination models on the grounds that these models assume that human capital investment decisions of individuals from different groups are separable. The assumed separability has the somewhat implausible — from a positive political economy perspective — implication that the privileged group has no incentive to preserve negative stereotypes, or equivalently, has no incentive to resist public efforts to eliminate negative stereotypes. To address this shortcoming, Moro and Norman (2003b) adopt a more general production technology under which they obtain cross-group effects operating through wages and job assignments in general equilibrium. In our case cross-group effects arise through human capital spillovers in neighborhoods (or more broadly networks and other social settings). And as a consequence, the advantaged group may have an incentive to resist integration.[5]

The remainder of the paper is organized as follows. In the next section, we lay out a benchmark model of statistical discrimination. In Section 3 we incorporate neighborhood effects into this benchmark model and provide a preliminary analysis of statistical discrimination in the presence of neighborhood effects, first under complete segregation, and then, under partial integration. We also present a couple of numerical examples that illustrate how, within this set-up, integration on a large-enough scale can qualitatively change the set of attainable equilibria. Sections 4 and 5

[5] Benabou (1993) combines the local spillovers that we focus on in our paper, with the global interactions resulting from production complementarities that are the focus of Moro and Norman (2003b), but does so in a framework where individual human capital is perfectly observable.

contain our main results, presented more formally. In Section 4 we discuss the conditions under which, starting from a discriminatory equilibrium, integration on a large enough scale can eliminate negative stereotypes. Section 5 follows with a discussion of the conditions under which integration and the consequent elimination of statistical discrimination is likely to be welfare-improving, and Section 6 concludes.

2 A benchmark model of statistical discrimination

2.1 The basic setup

Consider the following variant of the Coate and Loury (1993) model. There are two groups of workers, 1 and 2, and group membership is costlessly observable. Workers may pay a cost to become qualified, or may remain unqualified. This cost should be interpreted as being a net cost, taking into account any intrinsic benefits that may be derived from education. The distribution of costs within each group is given by the distribution function $F(c)$, which denotes the proportion of workers with cost below c. We allow for the possibility that $F(0) > 0$, namely that there exist some individuals for whom the intrinsic benefit from acquiring human capital exceeds the cost of doing so. The cost distributions are the same for each group.[6]

Firms may assign workers to one of two jobs. One of these may be done equally well by qualified and unqualified workers, and results in zero payoffs to both firms and workers. The other yields a positive aggregate payoff x_q if a qualified worker is assigned, and a negative payoff x_u otherwise. Nothing essential is lost by setting $x_q = 1$ and $x_u = -1$.

Let s_i be the proportion of workers in group i who choose to become qualified. Firms can observe these population proportions but cannot observe individual characteristics of workers prior to assignment. They observe a noisy signal which can take one of two values: P (positive) or N (negative). The probability that this signal is P when the worker is qualified is p, and the probability that this signal is P when the worker is unqualified is $q < p$.

When a firm observes a P worker belonging to group i, the posterior probability that the worker is qualified is

$$\theta(s_i) = \frac{ps_i}{ps_i + q\,(1 - s_i)} \tag{1}$$

Similarly, when a firm observes an N worker belonging to group i, the posterior probability that

[6] The assumption of identical cost distributions is relaxed when we consider neighborhood effects below.

the worker is qualified is

$$\varphi(s_i) = \frac{(1-p)\, s_i}{(1-p)\, s_i + (1-q)\,(1-s_i)} \tag{2}$$

Since $p > q$, $\theta(s_i) > \varphi(s_i)$ for all $s_i \in (0, 1)$. Firms will assign workers to the skilled task if the expected payoff is positive. For a P worker, this requires

$$2\theta(s_i) - 1 > 0$$

Similarly, for N workers to be allocated to the skilled tasks

$$2\varphi(s_i) - 1 > 0$$

Assume that workers get the entire surplus thus generated (this requires that P workers get paid more than N workers if both are assigned to the skilled task). Hence wages are

$$\begin{aligned} w_p(s_i) &= \max\{2\theta(s_i) - 1, 0\} \\ w_n(s_i) &= \max\{2\varphi(s_i) - 1, 0\} \end{aligned}$$

Here it is being assumed that wages can depend on worker type as well as the signal, which would show up as discrimination in a standard wage regression. But in equilibrium (see below) it will never be the case that workers with the same test result and job assignment will receive different wages. Firms will assign workers to jobs on the basis of the population composition.

A worker will wish to become skilled if doing so would yield an expected increase in the wage that exceeds the cost of becoming skilled. Hence a worker with cost c in group i will derive a positive net benefit from investing if and only if

$$pw_p(s_i) + (1-p)w_n(s_i) - c > qw_p(s_i) + (1-q)w_n(s_i)$$

which simplifies to yield

$$(p-q)\,(w_p(s_i) - w_n(s_i)) > c.$$

Let $b(s_i) = (p-q)\,(w_p(s_i) - w_n(s_i))$ denote the expected benefits of investing when the skill share in group i is s_i.

In any group with skill share s_i, assuming that the workers who are skilled have lower costs than those who are not, the cost of becoming skilled for the marginal worker is given by $F^{-1}(s_i)$. This marginal worker will wish to become skilled if and only if the benefits $b(s_i)$ from doing so exceed this cost. Equilibria correspond to states s_i for which $b(s_i) = c(s_i)$, where $c(s_i) = F^{-1}(s_i)$ is the cost of becoming skilled for the 'marginal' worker.

It can be shown that there exist three ranges for the skill share s_i such that (i) all workers are assigned to the low-skill task, regardless of signal, for s_i in the lowest range (ii) workers are assigned to the high-skill job if and only if they have a positive signal value for s_i in the intermediate range, and (iii) all workers are assigned to the high-skill task, regardless of signal, for s_i in the highest range. In the last case, workers with a positive signal value earn more than those with a negative signal value. These considerations imply that $b(s)$ is single-peaked, takes its maximum at some interior value of s, and $b(s) = 0$ for s sufficiently small. These general properties are depicted in Figure 1, which also depicts a cost function $c(s)$.[7]

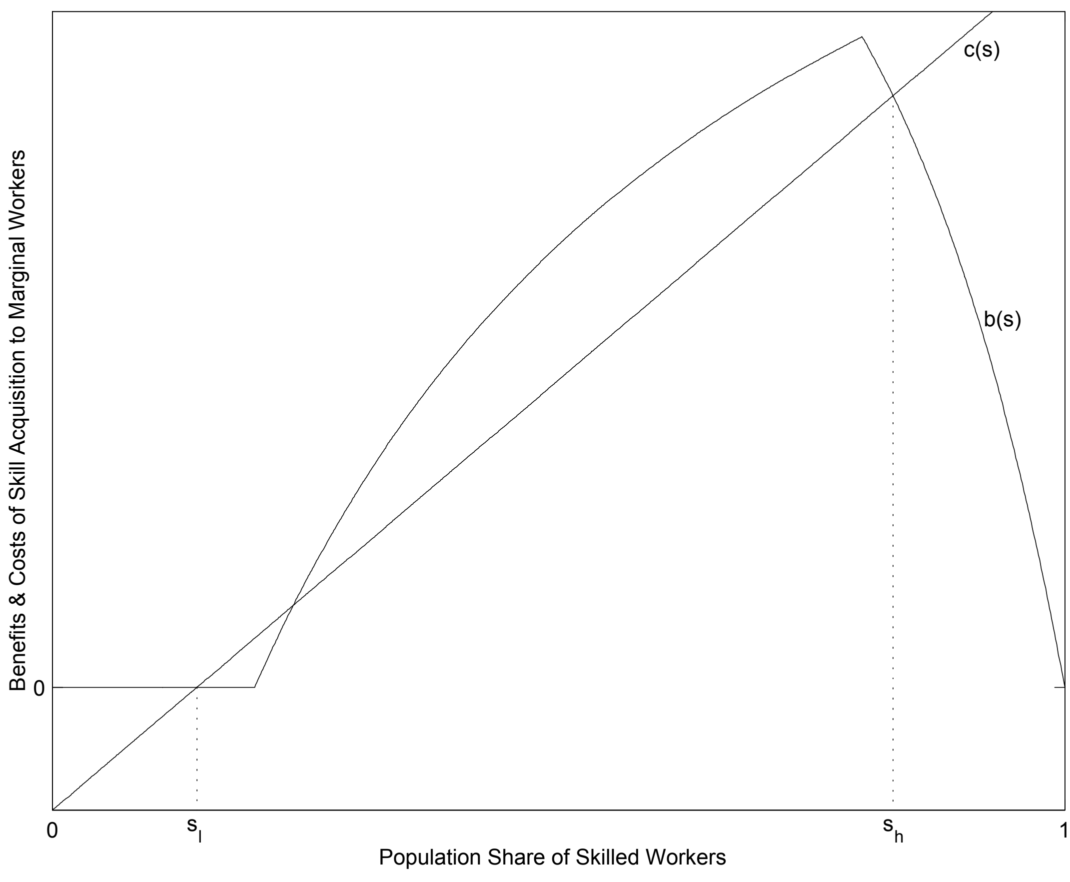

Figure 1. An Equilibrium with Negative Stereotypes

Consider first the case of symmetric stable equilibria $s_1 = s_2$ (in which both groups acquire identical levels of human capital and receive equal labor market treatment). If $F(0)$ is sufficiently small (so that relatively few individuals derive intrinsic benefits from becoming qualified), then

[7]The figure is based on the following specification: $p = 0.8$, $q = 0.2$, and costs are distributed uniformly with support $[-\varepsilon, \gamma - \varepsilon]$, where $\varepsilon = 0.05$ and $\gamma = 0.6$.

there will exist a stable equilibrium in which all workers are assigned to the low-skill task. There may, however, exist other symmetric equilibria. In the example of Figure 1, there are symmetric equilibria at (s_l, s_l) and (s_h, s_h).[8]

When multiple symmetric equilibria exist, there also exist asymmetric equilibria in which members of different groups receive different labor market treatment and make different human capital decisions. In the example of Figure 1, an equilibrium (s_l, s_h) with negative stereotypes about workers in group 1 exists. Here ex-ante identical workers end up with different levels of human capital. This captures an essential feature of statistical discrimination and its potential consequences when human capital choices are endogenous.

2.2 Stability

When multiple equilibria exist (as in Figure 1), not all of them will be dynamically stable. In order to identify stable equilibria, we need to specify the dynamics of skill shares under conditions of disequilibrium. It is natural to assume that at any skill profile (s_1, s_2), the skill share in group i will rise if and only if the benefits of skill acquisition exceed the costs for the marginal worker. We adopt the following simple specification:[9]

$$\dot{s}_i = b(s_i) - c(s_i).$$

Since human capital decisions, once made, are seldom reversible, the disequilibrium dynamics of skill shares should be interpreted as arising from the exit of older workers and the entry of younger ones into the labor market, as in Akerlof (1976).

In the absence of neighborhood effects the dynamics of skill shares in the two populations are independent and the stability of any equilibrium skill profile (s_1, s_2) requires only that for each s_i, the slope of the cost function exceed that of the benefit function at the equilibrium. In the example of Figure 1, any skill profile in which $s_i \in \{s_l, s_h\}$ for each i is a stable equilibrium. Hence there are four stable equilibria, of which two involve statistical discrimination.

In the presence of neighborhood effects the stability of equilibria is less easy to determine, since the cost distribution in each group depends on the skill acquisition in the other group as well as the level of neighborhood integration. This case is treated in Section 3 below.

[8] There is also a third symmetric equilibrium with intermediate levels of skill acquisition but this will be unstable in a sense to be made precise below.

[9] Any specification $\dot{s}_i = f(b(s_i) - c(s_i))$ with $f(0) = 0$, and f strictly increasing and Lipschitz continuous would leave our results intact.

2.3 The linear cost case

Conditions for the existence of multiple equilibria are easily obtained in the special case of a linear (marginal) cost function. Suppose costs are distributed uniformly on the interval $[-\varepsilon, \gamma - \varepsilon]$, where $0 < \varepsilon < \gamma$. In this case the distribution function is $F(c) = (c + \varepsilon)/\gamma$ and if the population share of workers who acquire skills is s, then the cost of skill acquisition for the marginal worker is $c(s) = \gamma s - \varepsilon$.

Define s_x as the threshold level of skill acquisition at which firms are indifferent between placing a worker with a positive signal in a skilled job versus an unskilled one. Specifically s_x is defined by the solution to $2\theta(s) - 1 = 0$, which yields

$$s_x = \frac{q}{p+q}.$$

If $c(s_x) > 0$, there will exist an equilibrium in which a share ε/γ of workers acquire skills and all workers are placed in unskilled jobs. The workers who acquire skills are precisely those whose costs of doing so are negative. We shall refer to this as the *low-skill equilibrium.* Let

$$\gamma_x = \frac{\varepsilon}{s_x}$$

denote the threshold value of the cost function parameter at which $c(s_x) = 0$. Then a low-skill equilibrium will exist whenever $\gamma > \gamma_x$. It is easy to verify that there can be at most one such equilibrium and that it must be stable.

We shall say that a *high-skill equilibrium* exists if there is some s such that $b(s) = c(s) > 0$. Define s_y as the level of skill acquisition at which $(b(s) + \varepsilon)/s$ is maximized. It is easily verified that there exists a unique s_y and that $\theta(s_y) > \frac{1}{2} > \varphi(s_y)$. In other words, when a share s_y of the population acquire skills, workers are placed in skilled jobs if and only if their signals are positive. Hence $b(s_y) = (p - q)\, w_p(s_y)$, and

$$s_y = \arg\max \frac{(p-q)\,(2\theta(s) - 1) + \varepsilon}{s}.$$

Define γ_y as the cost parameter at which $b(s_y) = c(s_y)$. Specifically, $\gamma_y = (b(s_y) + \varepsilon)/s_y$. Then a stable high-skill equilibrium exists if and only if $\gamma < \gamma_y$. The number of high-skill equilibria must be either zero or two, and in the case of the latter, exactly one of them will be stable. Although it is possible to get closed-form solutions for s_y and γ_y in terms of the parameters p, q, and ε, these expressions are cumbersome and are omitted here. Figure 2 illustrates geometrically the determination of the threshold skill shares s_x and s_y and the corresponding threshold cost functions.

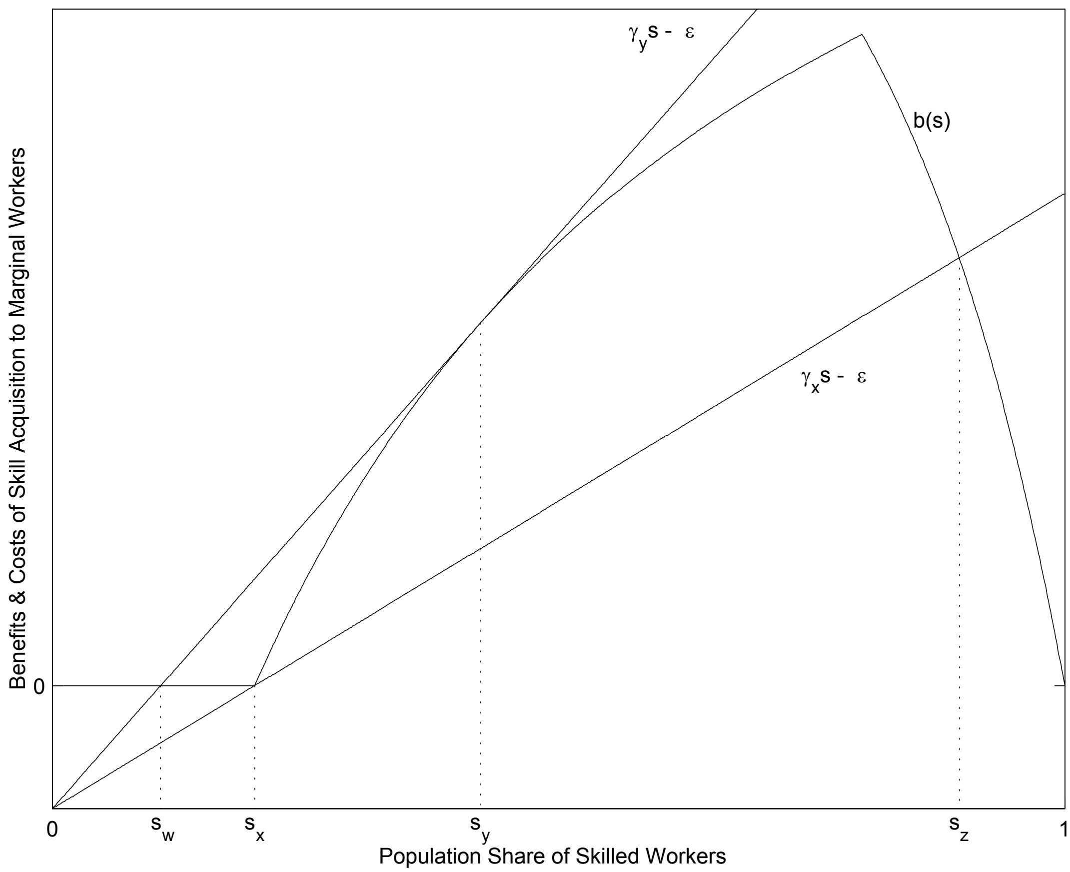

Figure 2. Cost Distributions Consistent with Statistical Discrimination

As is evident from the figure, an equilibrium with statistical discrimination exists if and only if $\gamma \in (\gamma_x, \gamma_y)$. The costs of skill acquisition must be neither too great nor too small. If they are too small, then there will be a unique symmetric high-skill equilibrium. If they are too large, there will be a unique symmetric low-skill equilibrium. In an intermediate range, multiple symmetric equilibria and hence also asymmetric equilibria can arise.

For future reference, we define s_w and s_z as follows. Let $s_z > s_x$ be the larger solution to $s\gamma_x - \varepsilon = b(s)$, and let $s_w < s_x$ be the smaller solution to $s\gamma_y - \varepsilon = b(s)$, as shown in Figure 2. When there exists a stable equilibrium (s_l, s_h) with statistical discrimination in this baseline model, it must be the case that $s_l \in [s_w, s_x]$ and $s_h \in [s_y, s_z]$. Once we allow for neighborhood effects, however, asymmetric equilibria can lie outside this range.

3 Incorporating neighborhood effects

The model above is based on the assumption, standard in the statistical discrimination literature, that the *ex-ante* costs of acquiring human capital are exogenously given and are the same for both groups. But if an asymmetric equilibrium is selected, then this assumption itself becomes difficult to sustain. Costs of becoming skilled will generally depend on the extent to which one's family members and neighbors have made prior investments in human capital. This in turn will depend on the extent of intergroup contact.

To address this, let β denote the share of the population belonging to group 1. The mean skill share in the population is then $\bar{s} = \beta s_1 + (1-\beta)s_2$. Let σ_i be the mean neighborhood skill level for individuals in group i (the share of their neighbors who become skilled).[10] When the groups are completely segregated and have no contact with each other, the mean skill share experienced by each group will simply be given by $\sigma_i = s_i$. On the other hand, under complete integration, each group would experience the same neighborhood skill share, which in turn would equal that in the population as a whole: $\sigma_1 = \sigma_2 = \bar{s}$. These are the two extremes within which neighborhood skill shares will lie for intermediate levels of segregation. We use the parameter $d \in [0,1]$ to represent the extent of segregation and specify

$$\sigma_i(s) = ds_i + (1-d)\,\bar{s}.$$

Lower values of d correspond to greater integration and hence a smaller distance between the neighborhood skill shares experienced by members of the two groups.[11]

We now generalize the determination of costs to allow for neighborhood effects. Let $G(c,\sigma)$ denote the proportion of individuals with costs below c in a group with mean neighborhood skill level σ. Positive spillovers from human capital accumulation at the neighborhood level are reflected in the assumption that G is increasing in σ. Although the costs of human capital accumulation may now differ across groups which experience different levels of neighborhood quality, it is assumed that the function G is identical across groups. Hence groups are not assumed to be innately different

[10] We use the term "neighborhood" loosely to indicate the set of people that individuals from group i associate with. Since residential proximity is one of several possible axes along which individuals interact, this may, but need not, coincide with a literal interpretation of neighborhoods in spatial terms.

[11] The parameter d is closely related to the index of dissimilarity commonly used in empirical studies of segregation. The index ranges between 0 and 1, and measures departures from a perfectly even distribution of the population across neighborhoods.

in any sense. Let $c(s, \sigma)$ denote the cost of becoming skilled for the marginal worker, when this worker experiences neighborhood quality σ and belongs to a group with skill share s.

3.1 Complete segregation

Neighborhood effects make statistical discrimination more likely to occur if the groups are segregated, since human capital spillovers are absent. Under complete segregation, each of the two groups can be analyzed independently. As before, suppose that costs are uniformly distributed on the interval $[-\varepsilon, \gamma - \varepsilon]$, but with variable support $\gamma = \gamma(\sigma)$. Higher neighborhood quality results in lower costs of human capital accumulation, so $\gamma(\sigma)$ is assumed to be strictly decreasing. The implied cost function is $c(s, \sigma) = s\gamma(\sigma) - \varepsilon$. Under complete segregation, $\sigma_i = s_i$ for $i = 1, 2$, so equilibrium requires that for each group, $c(s_i, s_i) = b(s_i)$. Even though the cost function is linear for any given level of neighborhood quality, the function $c(s, s)$ will be nonlinear in s, reflecting the fact that changes in neighborhood quality affect the distribution of costs.

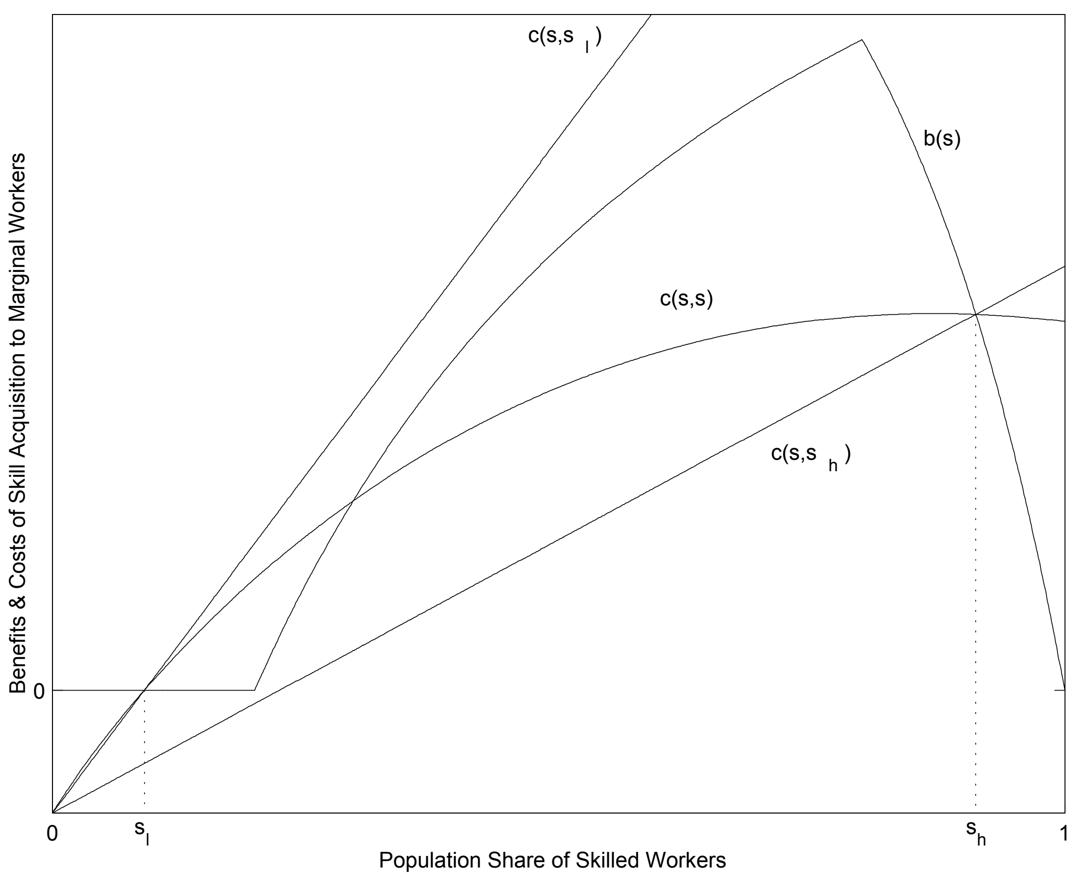

Figure 3. Statistical Discrimination under Complete Segregation

An equilibrium (s_l, s_h) with statistical discrimination under complete segregation is shown in Figure 3. Sufficient conditions for statistical discrimination under complete segregation are $c(s_x, s_x) > 0$ and $c(s_y, s_y) < b(s_y)$, where s_x and s_y are as defined in Section 2.3 (see also Figure 2). The first of these implies $s_x\gamma(s_x) - \varepsilon > 0$ or $\gamma(s_x) > \varepsilon/s_x = \gamma_x$. The second implies $b(s_y) > s_y\gamma(s_y) - \varepsilon$ or $\gamma(s_y) < (b(s_y) + \varepsilon)/s_y = \gamma_y$. Hence with neighborhood effects, a stable equilibrium with statistical discrimination exists under complete segregation if $\gamma(s_x) > \gamma_x$ and $\gamma(s_y) < \gamma_y$. Note that these conditions are not necessary. As before, if costs are neither too high nor too low in the relevant range, multiple equilibria will exist under complete segregation.

If a stable equilibrium with statistical discrimination exists under complete segregation, then there must exist *multiple* stable symmetric equilibria (with no statistical discrimination) at *all* levels of segregation.[12]

Proposition 1. *The skill profile (s_l, s_h) is an equilibrium with $d = 1$ if and only if (s_l, s_l) and (s_h, s_h) are both equilibria for all $d \in [0, 1]$.*

Hence there exists an equilibrium with statistical discrimination under complete segregation if and only if there exist multiple equilibria without statistical discrimination at all levels of segregation. We shall assume in the remainder of the paper that there exists an equilibrium with statistical discrimination under complete segregation and hence that there exist (at least) two symmetric equilibria (s_l, s_l) and (s_h, s_h) at all levels of segregation. This does not, of course, imply that there exists an asymmetric equilibrium at *all* levels of segregation. As we show below, there may exist a range of segregation levels at which no asymmetric equilibria exist, even in the presence of multiple symmetric equilibria.

3.2 Partial integration

When some degree of integration exists, the dynamics of skill shares are no longer independent in the two populations, since we have

$$\dot{s}_i = b(s_i) - c(s_i, \sigma_i) = b(s_i) - c(s_i, ds_i + (1-d)(\beta s_1 + (1-\beta)s_2))$$

so $\dot{s}_i$ depends on both s_1 and s_2.

Any curve in (s_1, s_2) space along which $\dot{s}_1 = 0$ is an *isocline* corresponding to the state variable s_1. Similarly the isoclines corresponding to the state variable s_2 are curves in (s_1, s_2) space along

[12] Proofs of all formal results are collected in the Appendix

which $\dot{s}_2 = 0$. All isoclines for a particular numerical example with partial integration are shown in Figure 4. The three flatter curves correspond to $\dot{s}_2 = 0$, while the three steeper ones represent all points at which $\dot{s}_1 = 0$. Any intersection of two isoclines corresponds to an equilibrium of the model (and a rest point of the dynamics). The arrows describe the direction in which s_1 and s_2 change in each of the sixteen regions of the state space. As in the example of Figure 1, there are four stable equilibria (represented by bullets) of which two involve negative stereotypes. The remaining five equilibria are unstable.

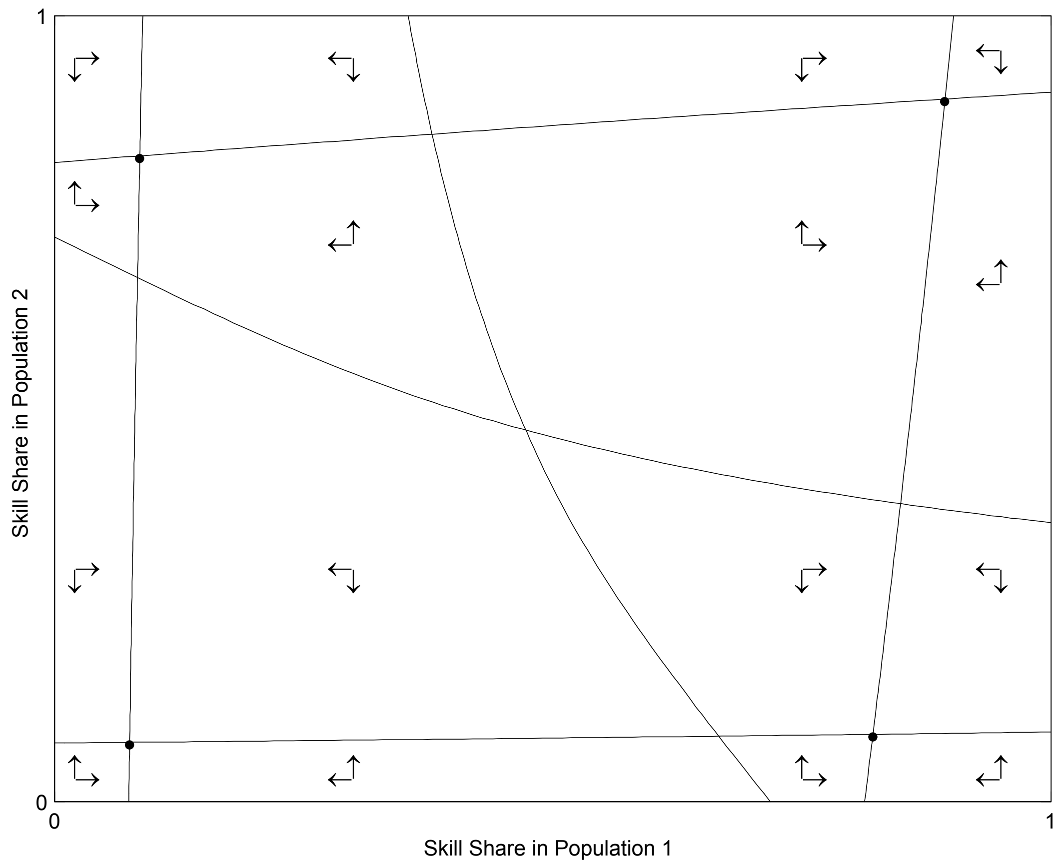

Figure 4. Isoclines and Stable Equilibria with Neighborhood Effects.

Of the four stable equilibria depicted in Figure 4, two are symmetric and two asymmetric. Of the latter, there is one in which the first group is disadvantaged and one in which it is advantaged. When speaking of equilibria with statistical discrimination in the remainder of this paper, we shall assume that the first group is disadvantaged, and denote the equilibrium (s_l, s_h) with the understanding that $s_l < s_h$.[13] Furthermore, we shall assume that at most one stable equilibrium

[13]This is without loss of generality because equilibria are invariant to a relabeling of groups provided that the

with this property exists.[14]

3.3 A first look at the effects of integration

Starting from a state (s_l, s_h) in which statistical equilibrium prevails, the effects of increasing integration on the human capital of the two groups will depend on whether or not integration leaves intact the *qualitative* properties of the set of equilibria. Suppose, for instance, that the economy is initially at the asymmetric equilibrium in the northwest corner of Figure 4, where the population 1 has lower skill levels. A decline in segregation, holding constant the two skill shares, will result in improved neighborhood quality for the disadvantaged group and worsened neighborhood quality for the advantaged group. At the original skill shares the result would be $c(s_1, \sigma_1) < b(s_1)$ and $c(s_2, \sigma_2) > b(s_2)$, implying $\dot{s}_1 > 0 > \dot{s}_2$. Assuming that the change in integration leaves intact the qualitative properties of the system (so that there remains an asymmetric equilibrium in which the first group is disadvantaged), the skill share will rise in the first group and fall in the second. This can be stated formally as follows.

Proposition 2. *Consider all values of d such that an asymmetric equilibrium (s_l, s_h) exists. Then s_l is strictly increasing in d, while s_h is strictly decreasing in d.*

Hence increasing integration, *if it allows negative stereotypes to persist*, will raise human capital levels in the disadvantaged group and lower them in the advantaged group. In this case small increases in integration benefit one group at the expense of the other but do so continuously. It is possible, however, that there exists a level of integration above which negative stereotypes are simply inconsistent with equilibrium. As integration proceeds beyond this critical *bifurcation* value, there is a *qualitative* change in the equilibrium properties of the model. In this case a small change in integration can result in a large and discontinuous shift in equilibrium skill levels. Since only symmetric equilibria remain, the economy must converge to a state in which both groups receive equal treatment. This leaves open the question of which of the two symmetric equilibria are reached.

An example is shown in Figure 5, which shows the phase diagrams corresponding to four different segregation levels. In this example, integration eliminates negative stereotypes but results in a shift

population share β is replaced with $1 - \beta$.

[14] The assumption that at most one stable equilibrium with statistical discrimination against the first group exists is consistent with as many as nine equilibria in the model as a whole (as in Figure 4), and is therefore quite unrestrictive. It is possible to obtain (but cumbersome to express) sufficient conditions on the primitives of the model that would guarantee this.

to the *less efficient* symmetric equilibrium. Hence the skill shares of *both* groups decline relative to the initial state in which negative stereotypes exist. Integration in this case raises the human capital of the disadvantaged group only until the bifurcation point, after which the skill shares of both groups collapse to levels below those at any asymmetric equilibrium. Equal treatment comes at a welfare cost to both groups, although the loss inflicted on the previously advantaged group is clearly greater.

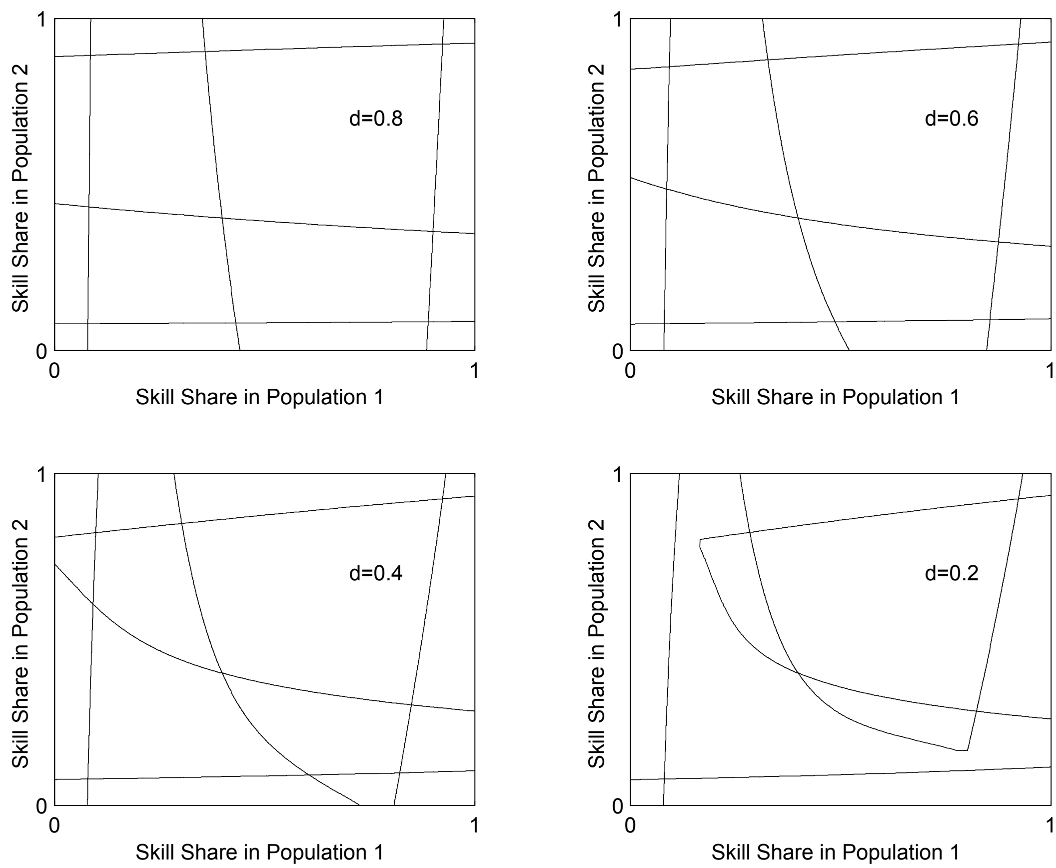

Figure 5. Integration Eliminates Negative Stereotpes but Lowers Skill Levels.

Integration can also eliminate negative stereotypes and result in *increases* in the human capital and welfare of both groups. An example of this is depicted in Figure 6. As before, there is a bifurcation point at which the equilibrium properties of the system undergo a qualitative change, and negative stereotypes become inconsistent with equilibrium. But, unlike the previous example, this time the economy is pushed into the basin of attraction of the symmetric equilibrium in which both groups have skill shares s_h. In this case, as long as integration proceeds beyond the critical threshold, it raises the skill levels of both groups relative to the initial state with negative stereotypes. Even the group which is advantaged under statistical discrimination can benefit from integration under such circumstances.

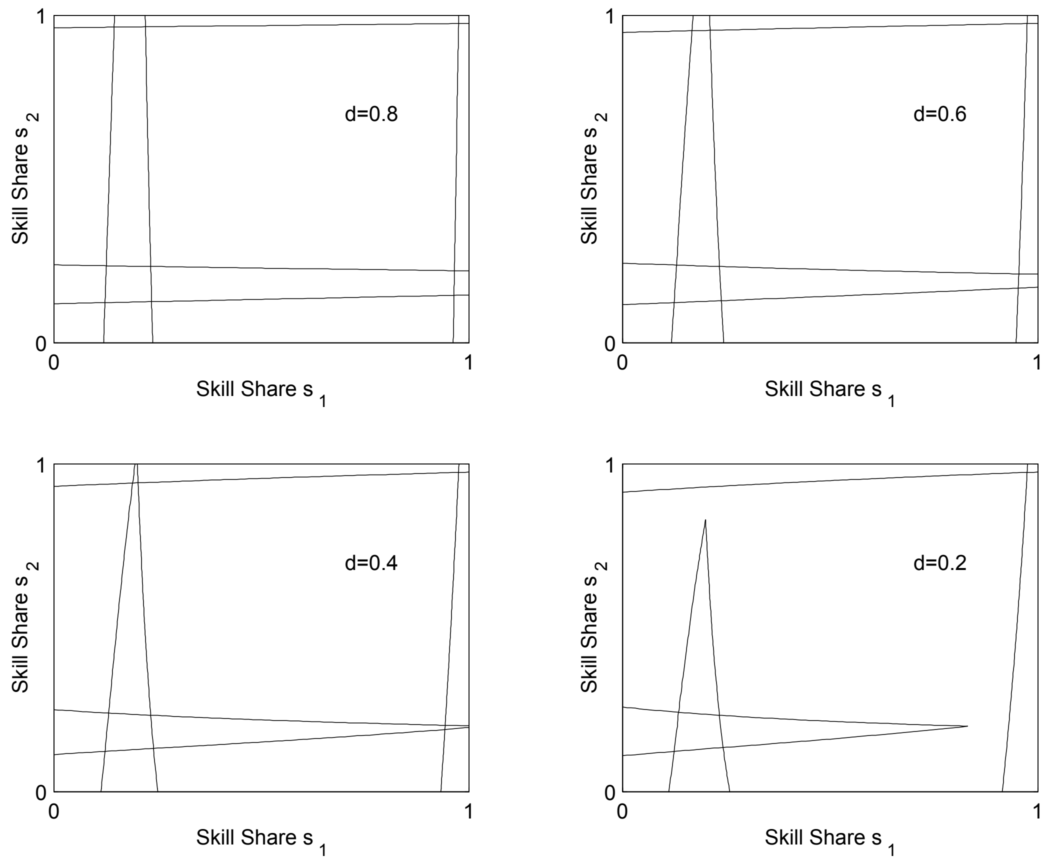

Figure 6. Integration Eliminates Negative Stereotypes and Raises Skill Levels.

These two examples illustrate not only that integration can eliminate negative stereotypes, but that its efficiency implications may depend in critical ways on underlying parameters. In the next section we identify conditions under which the elimination of negative stereotypes is efficiency-enhancing, and those in which it is efficiency-reducing. It turns out that the population share of the disadvantaged group, as well as the strength and curvature of neighborhood effects are critical in this regard.

4 Can integration eliminate negative stereotypes?

We shall say that *integration eliminates negative stereotypes* if there exists $\tilde{d} \in (0, 1)$ such that asymmetric equilibria exist if and only if $d \geq \tilde{d}$. If there exists no such $\tilde{d}$, then either discrimination can persist even under complete integration or cannot arise under any level of segregation. If integration eliminates negative stereotypes it triggers a shift from an asymmetric equilibrium to a symmetric one. As noted above, this shift can reduce skill shares in both population, or raise them in both populations.

Whether or not integration can eliminate negative stereotypes depends on both the population share of the disadvantaged group as well as the curvature of the relationship between neighborhood quality and the costs of human capital acquisition. In order to identify these effects analytically, we adopt a simple specification for costs. As in Section 2.2, suppose that costs are uniformly distributed with support $[-\varepsilon, \gamma - \varepsilon]$, so that the distribution function is $F(c) = (c + \varepsilon)/\gamma$. If a share s of workers acquire skills, then the cost of skill acquisition for the marginal worker is $c(s) = \gamma s - \varepsilon$. The parameter γ depends on neighborhood quality σ, and positive spillovers from human capital accumulation imply that $\gamma(\sigma)$ is strictly decreasing. We assume that the slope γ of the cost function depends on the neighborhood skill share σ as follows

$$\gamma(\sigma) = \bar{\gamma} - \sigma^a.$$

The parameter a, assumed to be strictly positive, captures the curvature of the relationship between σ and γ. When $a = 1$ the relationship is linear and the manner in which neighborhood quality affects costs is not sensitive to the initial neighborhood skill level. When $a > 1$, increases in neighborhood skill level result have their sharpest effect on costs when the initial skill level is already high. When the initial neighborhood skill level is low, improvements in neighborhood quality have negligible effects on the cost of human capital acquisition. When $a < 1$ on the other hand, the opposite is true: improvements in neighborhood quality have their sharpest effect on costs when the initial neighborhood quality is low. In communities with initially high skill levels, further improvements in neighborhood quality have a negligible impact on costs. Intuitively, integration will favor the disadvantaged community most (and hurt the advantaged group least) when the parameter a is small.

In general there may exist values of a such that equilibria with statistical discrimination cannot exist even under complete segregation. This can happen when a is sufficiently small (in which case both groups will have high levels of human capital in the unique symmetric equilibrium) or when a is sufficiently large (in which case both groups will have low levels of human capital in the unique symmetric equilibrium). Let $a_{\min}$ and $a_{\max}$ be defined, respectively, as the lower and upper bounds for a such that stable equilibria with statistical discrimination exist under complete segregation. In terms of the model's primitives, $a_{\min}$ is the supremum of the set of values of a such that $c(s,s) = b(s)$ has a unique solution $s > s_x$. Similarly $a_{\max}$ is the infimum of the set of values of a such that $c(s,s) = b(s)$ has a unique solution $s < s_x$. Since the cost function $c(s,s)$ is strictly increasing in a for all $s \in (0,1)$, the interval $[a_{\min}, a_{\max}] \subset [0,1]$ is uniquely defined. We shall assume that $a_{\min} < a_{\max}$, without which equilibrium statistical discrimination would be impossible for any parameter values.

By definition, for any $a \in (a_{\min}, a_{\max})$ there exists an asymmetric equilibrium (s_l, s_h) with $s_1 < s_2$ when $d = 1$. As noted in Proposition 1 above, s_l is increasing and s_h is decreasing in d. Integration, if it allows negative stereotypes to persist, will lead to greater skill accumulation in the disadvantaged group and less in the advantaged group. The result is therefore a narrowing of group inequality. It can further be shown that as long as negative stereotypes persist, both s_l and s_h are decreasing in β and a.[15] If the disadvantaged group's population share rises, both groups accumulate lower levels of human capital in equilibrium. This is a consequence of the fact that a rise in the population share of the disadvantaged group lowers neighborhood skill shares for both groups (as long as segregation is not complete). This raises the costs of skill acquisition and shifts the equilibrium in the intuitive direction. A rise in a also lowers skill shares of both groups because it reduces the benefits of contact with the advantaged group for the initially disadvantaged, while increasing the costs of contact with the disadvantaged group for the initially advantaged. Both effects result in lower equilibrium skill shares in each group.

Note that these comparative statics apply only if statistical discrimination persists in the face of parameter changes. As shown in the previous section, however, greater integration can make negative stereotypes unsustainable in equilibrium if it proceeds beyond some threshold. The following result establishes conditions under which this can occur

Theorem 1. *For any $\beta \in (0,1)$, there exist $a_m, a_n \in (a_{\min}, a_{\max})$ such that integration eliminates negative stereotypes for all $a \in (a_{\min}, a_m) \cup (a_n, a_{\max})$. Both a_m and a_n are decreasing in β.*

Hence integration can eliminate negative stereotypes for *any* population composition provided that the benefits of integration in low skill neighborhoods is sufficiently large relative to the costs in high skill neighborhoods, or alternatively, provided that the benefits of integration in low skill neighborhoods is sufficiently small relative to the costs in high skill neighborhoods. The threshold

[15] See the proof of Theorem 1 in the Appendix.

values $a_m(\beta)$ and $a_n(\beta)$ are both decreasing in β. The result is illustrated for a particular numerical specification of the model in Figure 7. The shaded regions of the figure represent the parameter ranges for which integration eliminates negative stereotypes.

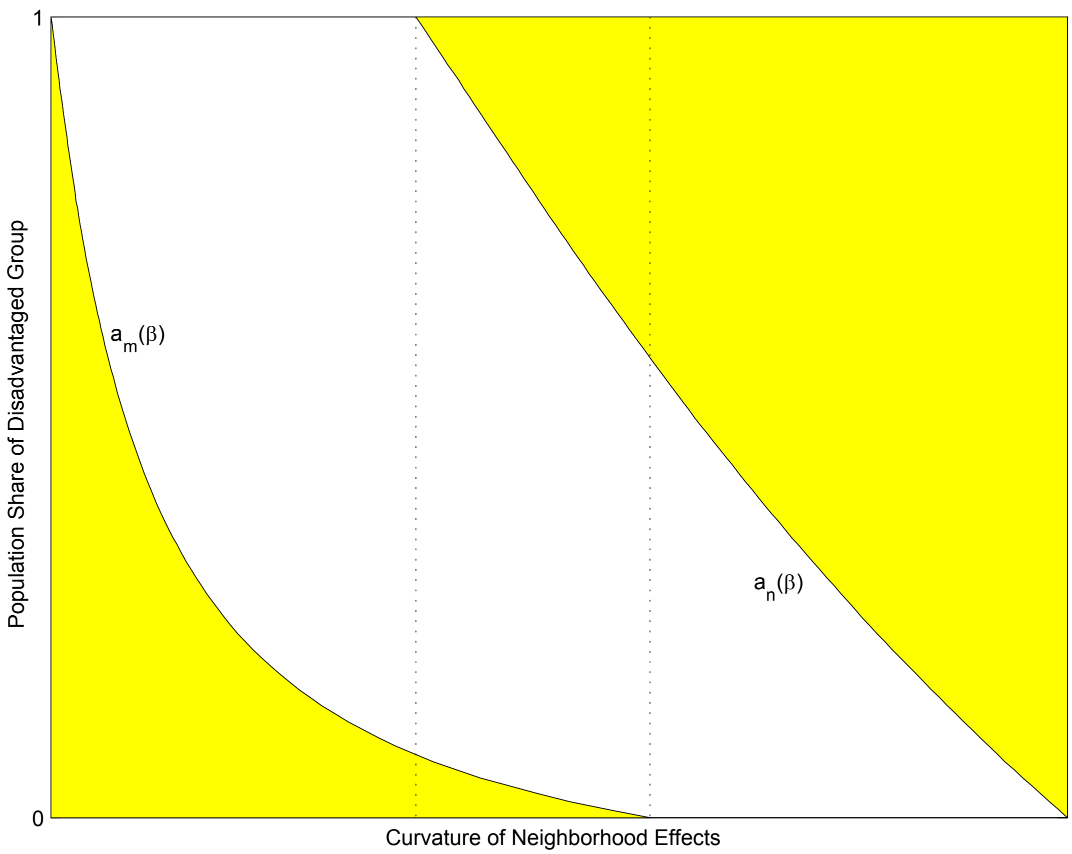

Figure 7. Parameter Range for which Integration Eliminates Negative Stereotypes

Theorem 1 implies that for any value of a between $a_m(1)$ and $a_n(0)$, there exists a range of values of β such that integration eliminates negative stereotypes. Specifically, the following is a Corollary of Theorem 1.

Corollary 1. *For any* $a \in (a_m(1), a_m(0))$ *there exists* $\beta_l > 0$ *such that integration eliminates negative stereotypes if* $\beta \in (0, \beta_l)$, *and for any* $a \in (a_n(1), a_n(0))$ *there exists* $\beta_h < 1$ *such that integration eliminates negative stereotypes if* $\beta \in (\beta_h, 1)$.

Hence integration can eliminate negative stereotypes if the population share of the disadvantaged group is sufficiently small, provided also that the benefits of integration to the disadvantaged are sufficiently large relative to the costs to the advantaged. It is also the case that integration can eliminate negative stereotypes if the population share of the disadvantaged group is sufficiently large, provided also that the benefits of integration to the disadvantaged are sufficiently small relative to the costs to the advantaged. Note that it is possible for $a_n(1)$ to be strictly smaller than $a_m(0)$, as in the example of Figure 7. In this case there will exists values of a (between the dashed lines in the figure) for which integration eliminates negative stereotypes if the population share of *either* group is sufficiently small.

5 Is integration welfare-enhancing?

From the perspective of horizontal equity, the elimination of negative stereotypes is, by definition, a good thing. And as we showed in the previous section, integration on a large enough scale can result in the elimination of negative stereotypes. Why then is there often such resistance to integration from members of the advantaged group? Our answer to this question is apparent in the firs of the two examples presented in the previous section. In the example depicted in Figure 5, integration on a large enough scale results in the elimination of negative stereotypes, and drives the economy to a symmetric equilibrium. The result, however, is that human capital levels and wages decline in both groups, and as a consequence, both groups are worse off than under the status quo with statistical discrimination.

Of course, as the example in Figure 6 demonstrates, the elimination of negative stereotypes as a consequence of large-scale integration can also result in a symmetric equilibrium where both groups are better off than under the status quo. In this section we explore the conditions under which each of these divergent outcomes is likely to occur. The population share, β, of the disadvantaged group, and the curvature of neighborhood effects, a, turn out to be critical.

We say that integration is welfare-enhancing if it eliminates negative stereotypes and results in an increase in skill levels in both populations relative to the status quo. It is welfare-reducing if it eliminates negative stereotypes and results in a decrease in skill levels in both populations.

Theorem 2. *For any $\beta \in (0,1)$, integration is welfare-enhancing if $a \in (a_{\min}, a_m)$, and welfare-reducing if $a \in (a_n, a_{\max})$.*

Here a_m and a_n are as defined in Theorem 1. Integration eliminates negative stereotypes and is welfare-enhancing if the benefits of intergroup contact to the disadvantaged group are sufficiently high relative to the costs of such contact to the advantaged group. Similarly, integration eliminates negative stereotypes and is welfare-reducing if the benefits of intergroup contact to the disadvantaged group are sufficiently small relative to the costs of such contact to the advantaged group. The result is quite intuitive. When a is small, intergroup contact raises skill shares in the stereotyped group more significantly than it lowers skill shares in the other group. Eventually a threshold is reached when the costs of human capital accumulation in the former group fall low enough to make self-fulfilling negative stereotypes unsustainable. At this point there is an increase in skill shares in *both* groups, although the increase is more rapid and significant in the previously stereotyped group. Eventually even the formerly advantaged group is better off, since they now experience positive human capital spillovers from all their neighbors, regardless of group membership. The reasoning is analogous for the case in which a is large, but the effects work in the opposite direction, lowering wages, skills and welfare in both groups.

According to Theorem 2, integration eliminates negative stereotypes and is welfare-enhancing if the parameters a and β lie in the lower-left shaded region in Figure 7. Similarly, integration eliminates negative stereotypes and is welfare-reducing if the parameters a and β lie in the upper-right shaded region in Figure 7. An immediate consequence of this is that when a is sufficiently small, integration is welfare-enhancing if β is also sufficiently small. When a is sufficiently large, on the other hand, integration is welfare-reducing if β is also sufficiently large. Specifically, the following is implied by Theorem 2.

Corollary 2. *For any $a \in (a_m(1), a_m(0))$ integration is welfare-enhancing if $\beta \in (0, \beta_l)$, and for any $a \in (a_n(1), a_n(0))$ integration is welfare-reducing if $\beta \in (\beta_h, 1)$.*

Here β_l, and β_h are as defined in Corollary 1. If the disadvantaged group is a sufficiently small minority, both groups can benefit from a policy of integration if negative stereotypes are eliminated as a result. In this case one might expect widespread popular support for integrationist policies even within the ranks of the advantaged group. On the other hand, integration may be welfare reducing for both groups if the initially disadvantaged group is a sufficiently large majority.

6 Conclusion

Neighborhood effects (or more broadly, a variety of local complementarities) have featured prominently in analyses of human capital formation. Somewhat surprisingly, such effects have been ignored in the literature on statistical discrimination in labor market settings. This is the case even in that strand of the literature that emphasizes how negative stereotypes might be perpetuated through the endogenous human capital acquisition decisions of those subject to such beliefs. We see the statistical discrimination literature and the literature on neighborhood effects in human capital accumulation as being deeply complementary, and have taken a step towards bridging the gap between the two.

We did so by introducing neighborhood effects into an otherwise standard model of statistical discrimination in job assignment. With the introduction of neighborhood effects, the level of segregation in the economy becomes salient in determining the likelihood and extent of statistical discrimination. We showed that starting from a situation where statistical discrimination is prevalent, integration on a large enough scale can, if neighborhood effects operate in a sufficiently asymmetric fashion, eliminate negative stereotypes by making it impossible for a stable discriminatory equilibrium to be sustained. Whether this results in a welfare improvement or results instead in both groups being worse off than under statistical discrimination depends on the population share of the initially disadvantaged group and the direction of the asymmetry in the impact of neighborhood effects. In metropolitan areas in which the stereotyped group is relatively small, and if human capital spillover effects are most powerful at low levels of skill accumulation, one would expect vigorous integrationist policies to be both uniformly welfare-enhancing in the long run, and to enjoy widespread popular support.

Appendix

Proof of Proposition 1. Suppose that (s_l, s_h) is an equilibrium with $d = 1$. Since $d = 1$, $\sigma_1 = s_l$ and $\sigma_2 = s_h$. Hence equilibrium implies $b(s_l) = c(s_l, s_l)$ and $b(s_h) = c(s_h, s_h)$. Now consider the symmetric skill allocation (s_l, s_l) for any level of $d \in [0, 1]$. At this allocation $\sigma_1 = \sigma_2 = s_l$ regardless of d. Since $b(s_l) = c(s_l, s_l)$, this must be an equilibrium. An identical argument shows that (s_h, s_h) must also be an equilibrium for all $d \in [0, 1]$. This proves that (s_l, s_h) is an equilibrium with $d = 1$ only if (s_l, s_l) and (s_h, s_h) are both equilibria for all $d \in [0, 1]$. To show that the converse is also true, note that if (s_l, s_l) and (s_h, s_h) are both equilibria for all $d \in [0, 1]$ then in particular, they are both equilibria for $d = 1$. This implies $b(s_l) = c(s_l, s_l)$ and $b(s_h) = c(s_h, s_h)$. Since $\sigma_1 = s_1$ and $\sigma_2 = s_2$ when $d = 1$, it follows that (s_l, s_h) is an equilibrium. ■

Proof of Proposition 2. Starting from any equilibrium (s_l, s_h), a decline in segregation will raise the neighborhood quality experienced by the disadvantaged group and raise the neighborhood quality experienced by the advantaged group at any state (s_1, s_2) that lies within a sufficiently small neighborhood of the initial equilibrium. This implies that for any s_2 in such a neighborhood, the value of s_1 at which $b(s_1) = c(s_1, \sigma_1)$ will be higher than it was prior to the decline in segregation. In other words, the fall in segregation will shift the relevant $\dot{s}_1 = 0$ isocline to the right. Analogous reasoning may be used to show that the relevant $\dot{s}_2 = 0$ isocline will shift down. If the two isoclines continue to intersect, there will be a new equilibrium close to the original one, with the following properties: (i) the original equilibrium will lie in the basin of attraction of the new one, and (ii) the dynamics of skill shares will satisfy $\dot{s}_2 < 0 < \dot{s}_1$ at the original equilibrium state. Hence a decline in segregation raises the equilibrium skill share of the disadvantaged group and lowers that of the advantaged group. ■

Proof of Theorem 1. First consider all values of a and β such that an asymmetric equilibrium (s_l, s_h) exists. We shall show that both s_l and s_h are decreasing in a and β. Starting from any equilibrium (s_l, s_h), an increase in a raises costs $c(s_i, \sigma_i)$ for each group at the original state. Hence at any state (s_1, s_2) that lies within a sufficiently small neighborhood of the initial equilibrium (s_l, s_h) we have $c(s_i, \sigma_i) > b(s_i)$ and hence $\dot{s}_i < 0$ for $i = 1, 2$. This implies that for any s_2 in such a neighborhood, the value of s_1 at which $b(s_1) = c(s_1, \sigma_1)$ will be lower than it was prior to the increase in a. In other words, the rise in a will shift the relevant $\dot{s}_1 = 0$ isocline to the left. Analogous reasoning may be used to demonstrate that the relevant $\dot{s}_2 = 0$ isocline will shift down. If the two isoclines continue to intersect, there will be a new equilibrium close to the original one, with the following properties: (i) the original equilibrium will lie in the basin of attraction of the new one, and (ii) the dynamics of skill shares will satisfy $\dot{s}_1 < 0$ and $\dot{s}_2 < 0$ at the original equilibrium state. Hence a rise in a shifts both equilibrium skill shares down. The effects of changes in β is similar: starting from any equilibrium (s_l, s_h), an increase in β lowers neighborhood quality for each group and hence raises costs $c(s_i, \sigma_i)$ for each group at the original state. This is exactly the effect of an increase in a, so the reasoning used for that case applies also here: a rise in β shifts both equilibrium skill shares down. Hence both s_l and s_h are decreasing in a and β.

Let $\bar{s}_{xz} = \beta s_x + (1-\beta)s_z$ denote the neighborhood skill share that would prevail under complete integration at an equilibrium in which $(s_l, s_h) = (s_x, s_z)$. Define a_m as the value of a at which $\gamma(\bar{s}_{xz}) = \gamma_x$. When $a = a_m$, $\gamma(s_z) < \gamma_x < \gamma(s_x)$ so there exists an asymmetric equilibrium under complete segregation. Hence $a_m \in (a_{\min}, a_{\max})$. Note that $\gamma(\bar{s}_{xz}) = \gamma_x$ implies that $c(s_x, \bar{s}_{xz}) = 0$ and $c(s_z, \bar{s}_{xz}) = b(s_z)$. This implies that when $a = a_m$, there is an asymmetric equilibrium $(s_l, s_h) = (s_x, s_z)$ under complete integration ($d = 0$). At any $a < a_m$ there is no asymmetric equilibrium under complete integration since, as shown above, s_l and s_h are both decreasing in a and $s_l \in [s_w, s_x]$ under complete integration. Hence integration eliminates negative stereotypes for all $a \in (a_{\min}, a_m)$. That a_m is decreasing in β follows from the definition of a_m and the fact that both s_l and s_h are decreasing in a and β.

Next let $\bar{s}_{wy} = \beta s_w + (1-\beta)s_y$ denote the neighborhood skill share that would prevail under complete integration at an equilibrium in which $(s_l, s_h) = (s_w, s_y)$. Define a_n as the value of a at which $\gamma(\bar{s}_{wy}) = \gamma_y$. When $a = a_n$ $\gamma(s_w) < \gamma_y < \gamma(s_y)$ so there exists an asymmetric equilibrium under complete segregation. Hence $a_n \in (a_{\min}, a_{\max})$. Note that $\gamma(\bar{s}_{wy}) = \gamma_y$ implies that $c(s_w, \bar{s}_{wy}) = 0$ and $c(s_y, \bar{s}_{wy}) = b(s_y)$. This implies that when $a = a_n$, there is an asymmetric equilibrium $(s_l, s_h) = (s_w, s_y)$ under complete integration ($d = 0$). At any $a > a_n$ there is no asymmetric equilibrium under complete integration since, from Proposition 1, s_l and s_h both decreasing in a and $s_l \in [s_w, s_x]$ under complete integration. Hence integration eliminates negative stereotypes for all $a \in (a_n, a_{\max})$. That a_n is decreasing in β follows from the definition of a_n and the fact that both s_l and s_h are decreasing in a and β. ■

Proof of Theorem 2. From the proof of Theorem 1, when $a = a_m$ and $d = 0$ there is an asymmetric equilibrium (s_x, s_z). Since the lower skill share is decreasing in a and d at asymmetric equilibria (from Proposition 1), for any $a < a_m$ there exists some segregation level $\tilde{d} \in (0, 1)$ such that an asymmetric equilibrium (s_l, s_h) exists with $s_l(a, \beta, \tilde{d}) = s_x$ and $s_h(a, \beta, \tilde{d}) > s_z$.

(This is the segregation level beyond which negative stereotypes are eliminated.) Let N_δ denote a neighborhood of the equilibrium point (s_l, s_h) such that the Euclidean distance between any point in N_δ and the equilibrium (s_l, s_h) is at most δ. We claim that if δ is sufficiently small, then at all points $(s_1, s_2) \in N_\delta$ at which $\dot{s}_1 = 0$, we must have $s_2 \leq s_h$. To see why, note that at any (s_1, s_2) at which $s_1 \leq s_l$ and $s_2 > s_h$, $c(s_1, \sigma_1) < b(s_1) = 0$. This implies that $\dot{s}_1 = 0$ with $s_2 > s_h$ only if $s_1 > s_l$. However, since $s_l = s_x$, $c(s_1, \sigma_1) < b(s_1)$ for all $s_1 > s_l$ if $s_2 > s_h$, $(s_1, s_2) \in N_\delta$, and δ is sufficiently small. This shows that the isocline $\dot{s}_1 = 0$ does not extend past the isocline $\dot{s}_2 = 0$ at their intersection (s_l, s_h). (This is the case depicted in Figure 6). Any decline in d past $\tilde{d}$ therefore leaves the economy in the basin of attraction of the symmetric equilibrium in which both groups have high levels of human capital.

Next consider the case $a > a_{n.}$ From the proof of Proposition 3, when $a = a_n$ and $d = 0$ there is an asymmetric equilibrium (s_w, s_y). Since the higher skill share is decreasing in a and increasing in d at asymmetric equilibria (from Proposition 1), for any $a > a_n$ there exists some segregation level $\tilde{d} \in (0, 1)$ such that an asymmetric equilibrium (s_l, s_h) exists with $s_h(a, \beta, \tilde{d}) = s_y$ and $s_l(a, \beta, \tilde{d}) < s_w$. (This is the segregation level beyond which negative stereotypes are eliminated.) Let N_δ denote a neighborhood of the equilibrium point (s_l, s_h) such that the Euclidean distance between any point in N_δ and the equilibrium (s_l, s_h) is at most δ. We claim that if δ is sufficiently small, then at all points $(s_1, s_2) \in N_\delta$ at which $\dot{s}_2 = 0$, we must have $s_1 \geq s_l$. To see why, note that at any (s_1, s_2) at which $s_2 \geq s_h$ and $s_1 < s_l$, $c(s_2, \sigma_2) > b(s_2)$. This implies that $\dot{s}_2 = 0$ with $s_1 < s_l$ only if $s_2 < s_h$. However, since $s_h = s_y$, $c(s_2, \sigma_2) > b(s_2)$ for all $s_2 < s_h$ if $s_1 < s_l$, $(s_1, s_2) \in N_\delta$, and δ is sufficiently small. This shows that the isocline $\dot{s}_2 = 0$ does not extend past the isocline $\dot{s}_1 = 0$ at their intersection at (s_l, s_h). (This is the case depicted in Figure 5). Any decline in d past $\tilde{d}$ therefore leaves the economy in the basin of attraction of the symmetric equilibrium in which both groups have low levels of human capital. ■

References

[1] Aigner, D.J. and G.C. Cain (1977). "Statistical Theories of Discrimination in the Labor Market," *Industrial and Labor Relations Review*, January, 30: 175-187.

[2] Akerlof, G. (1976). "The Economics of Caste and of the Rat Race and Other Woeful Tales." *Quarterly Journal of Economics* 90, 599–617.

[3] Arnott, R. and R. Rowse (1987). "Peer Group Effects and Educational Attainment," *Journal of Public Economics*, 32: 287-305.

[4] Arrow, K.J. (1973). "The Theory of Discrimination." In Ashenfelter and Rees, editors, *Discrimination in Labor Markets.*

[5] Benabou, R. (1993). "Workings of a City: Location, Education, and Production," *Quarterly Journal of Economics*, August, 108: 619-652.

[6] Benabou, R. (1996a). "Equity and Efficiency in Human Capital Investment: The Local Connection," *Review of Economic Studies*, April, 63: 237-264.

[7] Benabou, R. (1996b). "Heterogeneity, Stratification and Growth: Macroeconomic Implications of Community Structure and School Finance," *American Economic Review*, June, 86: 584-609.

[8] Borjas, G.J. (1995). "Ethnicity, Neighborhoods, and Human Capital Externalities," *American Economic Review*, 85: 365-390.

[9] Borjas, G.J. and M.S. Goldberg (1979). "Biased Screening and Discrimination in the Labor Market," *American Economic Review*, December, 68: 918-922.

[10] Coate, S. and G.C. Loury (1993). "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83: 1220–1240.

[11] Coleman, J. et al. (1966). *Equality of Educational Opportunity.* Washington: U.S. Department of Health, Education, and Welfare.

[12] De Bartolome, C. (1990). "Equilibrium and Inefficiency in a Community Model with Peer Group Effects," *Journal of Political Economy*, 98: 110-133.

[13] Henderson, V., P. Mieszkowski and Y. Sauvageau (1978). "Peer Group Effects and Educational Production Functions," *Journal of Public Economics*, 10: 97-106.

[14] Ioannides, Y. (2001). "Nonlinear Neighborhood Interactions and Intergenerational Transmission of Human Capital," *mimeo*, Department of Economics, Tufts University.

[15] Kremer, M. (1997). "How Much Does Sorting Increase Inequality?" *Quarterly Journal of Economics,* 112: 115-139.

[16] Loury, G.C. (2002). *The Anatomy of Racial Inequality.* Harvard University Press.

[17] Lucas, R., Jr. (1988). "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 22: 3-42.

[18] Lundberg, S.J. and R. Startz (1983). "Private Discrimination and Social Intervention in Competitive Labor Markets," *American Economic Review*, June, 73: 340-347.

[19] Moro, A. and P. Norman (2003a). "Affirmative Action in a Competitive Economy," *Journal of Public Economics* (forthcoming).

[20] Moro, A. and P. Norman (2003b). "A General Equilibrium Model of Statistical Discrimination." *Journal of Economic Theory*, forthcoming.

[21] Norman, P. (2003). "Statistical Discrimination and Efficiency," *Review of Economic Studies*, 70: 615-627.

[22] Phelps, E. (1972). "The Statistical Theory of Racism and Sexism." *American Economic Review* 62: 659–661.

[23] Schwab, S. (1986). "Is Statistical Discrimination Efficient?", *American Economic Review*, 76: 229-234.

[24] Sethi, R. and R. Somanathan (2002). "Inequality and Segregation," *mimeo*, Department of Economics, Barnard College, Columbia University.

[25] Summers, A. and B. Wolfe (1977), "Do Schools Make a Difference?," *American Economic Review*, 68: 639-652.