Published in final edited form as:

J Polit Econ. 2013 August ; 121(4): 747–774. doi:10.1086/671180.

Valuing Diversity

Roland G. Fryer Jr. and
Harvard University

Glenn C. Loury
Brown University

Abstract

This paper explores the economics of diversity-enhancing policies. A model is proposed in which heterogeneous agents, distinguished by skill level and social identity, purchase productive opportunities in a competitive market. We analyze policies designed to raise the status of a disadvantaged identity group. When agent identity is contractible, efficient policy grants preferred access to slots but offers no direct assistance for acquiring skills. When identity is not contractible, efficient policy provides universal subsidies to skill development when the fraction of the disadvantaged group at the skill development margin is larger than their share at the slot assignment margin.

This is the next and the more profound stage of the battle for civil rights. We seek not just freedom but opportunity. We seek not just legal equity but human ability, not just equality as a right and a theory but equality as a fact and equality as a result. (President Lyndon B. Johnson, Howard University, 1965)

I. Introduction

When choosing which students to admit, employees to hire, candidates to slate, or firms to patronize, the social identity of those selected—that is, an agent’s race, sex, age, nationality, religion, ethnicity, or caste—can be a matter of great importance. As a consequence, regulations intended to achieve more diversity in the ranks of the chosen—policies going under the rubric of “affirmative action,” “positive discrimination,” or (less neutrally) “reverse discrimination”—have been promulgated in many societies throughout the world. This paper examines the welfare economics of such diversity-promoting public regulation.

Consider a few examples. Where there are sharp sectarian divisions—Lebanon, Indonesia, Pakistan, Iraq—political stability can hinge on maintaining ethnic balance in the military ranks and on distributing coveted political offices so that no single group has disproportionate influence. In the United States, selective colleges and universities often feel obliged to take actions calculated to enhance the racial diversity of their student bodies. Amid rioting and civil unrest, France has designed policies to ensure more diversity in firms and selective institutions of higher education. Elsewhere in Europe, political parties have

mandated that female candidates be adequately represented on their electoral lists. To ensure that wealth is more widely distributed in post-apartheid South Africa, a policy of Broad-Based Black Economic Empowerment has been enacted that sets numerical standards of black representation that companies are encouraged to meet. In the wake of widespread ethnic rioting that erupted in Malaysia, a New Economic Policy was instituted in 1969 that created quotas and preferences for ethnic Malays in public contracting, employment, and education. In India, so-called scheduled castes and tribes enjoy preferred access to university seats and government jobs by constitutional mandate, though amid fierce controversy.¹ These varied programs differ in many details, but we refer here to all such diversity-enhancing efforts as “affirmative action.”²

Affirmative action policies generally entail the preferential treatment of persons who possess certain social identities based on a presumption that, on the average, those belonging to the preferred group are less effective in the competition for scarce positions because of some preexisting social handicap. Yet, exogenous between-group differences necessarily make diversity a costly commodity when the supply of opportunity is limited. The reality of unequal development between groups creates an unavoidable economic problem for those wishing to undertake affirmative action: enhanced access for a genuinely disadvantaged group to much sought-after productive opportunities cannot be achieved without altering selection standards, distorting human capital investment decisions, or both.³ So the relevant economic problem, which we study here, is to consider how these costs should be conceptualized and how they can be minimized.

Our analysis of the welfare economics of affirmative action policies is motivated by two thematic questions in particular:

1. Where in the economic life cycle should preferential treatment be most emphasized: before productivities have been essentially determined or after?
2. How do public policies that valorize a nonproductive trait—that is, identity—affect private incentives to become more productive?

To explore these questions, we develop a simple two-stage model. In the *ex ante* stage, heterogeneous agents, distinguished by their costs of skill acquisition and their social identities, decide whether to invest in skills. In the *ex post* stage, these agents, distinguished now by productivity and social identity, enter a competitive market in which they purchase “slots” (i.e., productive opportunities) that enable them to use their skills. One social identity group is presumed to be “disadvantaged” in the sense that they have higher costs of skill

¹Many other examples could be given, from countries such as the Philippines, Nigeria, and Sri Lanka. For a comprehensive review and assessment of these policies in global perspective, see Sowell (2004). On maintaining ethnic diversity in military selection, see Klitgaard (1986). On racial preferences at US colleges and universities, see Bowen and Bok (1998). On caste and ethnic preferences in India, see Galanter (1992) and Deshpande (2006).

²There has been much heated debate about the fairness of affirmative action. We do not take up that question here. While fairness issues are an important concern, these policies—and the controversy they inevitably inspire—can be found virtually everywhere. For this reason, our focus in this paper is on how greater racial or ethnic diversity in a hierarchy of positions can be efficiently achieved. For a discussion of the social justice issues raised by affirmative action and other racially egalitarian policies, see Loury (2002, chap. 4).

³At least this will be so in the short run when a group’s preexisting social handicaps cannot be easily or quickly ameliorated. See Fryer and Loury (2005a) for a detailed discussion of some of the usually overlooked, yet unavoidable, trade-offs associated with affirmative action policies.

acquisition on the average. With no policy intervention the disadvantaged will be less than proportionally represented among slot owners, compared to their presence in the population, since agents face a common slots price in market equilibrium.

This relative underrepresentation can give rise to a demand for some sort of diversity-enhancing measures. We consider policy interventions to enhance opportunity for agents in the disadvantaged group. Designing an efficient policy of this kind is posed as an elementary economics problem, and the simplicity of our framework allows us to easily derive a number of results. What makes our analysis novel, relative to the existing literature on affirmative action, is our focus on second-best efficiency questions and the attention we give to the *visibility* and the *timing* dimensions of affirmative action policies.⁴

The *timing* issue has to do with finding an ideal point in the developmental process to introduce a preference. We distinguish in the model between the *ex ante* and the *ex post* stages of production. Given that we assume that a productivity gap already exists, an *ex post* preference offers a competitive edge to less productive agents in the disadvantaged group. By contrast, an *ex ante* preference promotes the competitive success of the disadvantaged by fostering their prior acquisition of skills. That is, *ex ante* policies operate on the *development margin*, while *ex post* policies operate on the *assignment margin*.

The *visibility* dimension concerns an informational constraint one often encounters with affirmative action, reflected in the distinction we draw between *sighted* and *blind* policy environments. Under sightedness, assignment standards and development subsidies can be tailored to group membership at the individual level. Sighted policies are overtly discriminatory in that otherwise-similar agents from different groups are treated differently. Blind policies, in contrast, are tacitly discriminatory. They have their impact by placing a premium on some nonidentity traits that are known to be more prevalent in the preferred population. Though they are facially neutral in their treatment of groups, blind preferential policies have been intentionally chosen to have group-disparate effects.

Combining these distinctions of visibility and of timing generates a 2×2 conceptual matrix that captures the main contours of affirmative action as it is practiced in the real world: job reservations, contract set-asides, distinct admissions standards, and race-normed ability tests; all exemplify sighted *ex post* preferences. Instances of sighted *ex ante* preferences include minority scholarship funds, group-targeted skills development programs, and costly outreach and recruitment efforts that encourage an underrepresented group to prepare for future opportunities.

On the other hand, automatic admission for the top 10 percent of a state's high school classes, waiving a mandate that college applicants submit test scores, selecting among applicants partly by lot, or introducing nonidentity factors that are unrelated to performance into the evaluation process are all examples of blind *ex post* preferences.⁵ And, since there must be some group disparity in the distribution of endowments (otherwise, no policy

⁴Earlier papers on the economics of affirmative action policies include Welch (1976), Lundberg and Startz (1983), Coate and Loury (1993), Moro and Norman (2003), and Fryer and Loury (2005b). For a review of the evidence on the effectiveness of these policies, see Holzer and Neumark (2000). For a broad policy discussion, see Fryer and Loury (2005a).

promoting group equality would be needed), a blind ex ante preference can always be put in place by subsidizing for everyone those skill-enhancing actions from which agents in a preferred group can derive the most benefit.⁶

Within our simple framework, we produce two theoretical results on affirmative action policy. When group identity is visible and contractible, then, given some target representation rate in the slots market for the disadvantaged group, the problem is identical to solving the laissez-faire slot allocation problem for each group individually. This result also implies that the optimal ex post stage sighted intervention can be implemented by subsidizing the acquisition of slots by members of the disadvantaged group at a fixed rate. Because this identity-based subsidy raises the competitive price of slots at the ex post stage, it prices marginal members of the advantaged group out of the slots market. It also raises skill acquisition among the disadvantaged and lowers it among the advantaged—relative to laissez-faire—in a socially optimal fashion.

The second result pertains to the regime of blind affirmative action. We show two things: first (under a “convexity of the likelihood ratio” condition), the efficient blind ex post stage policy is a mix of competitive bidding and random assignment of slots; second, the efficient blind ex ante stage policy entails either a universal subsidy to skill enhancement or a universal tax on skill enhancement depending, respectively, on whether at the efficient blind allocation the disadvantaged group is better represented on the development margin or on the assignment margin.

The paper proceeds as follows. Section II introduces a simple model of production with investment in skills and competition for positions in a hierarchy. Section III brings affirmative action policy into the model to formally represent public interventions that expand opportunity for a disadvantaged group. Section IV presents concluding remarks.

II. The Model

A. Basic Building Blocks

Imagine a world in which agents purchase “slots” to produce widgets. These agents form a continuum of unit measure and belong to one of two social identity groups, $i \in \{a, b\}$. The fraction of agents in group i is denoted $\lambda_i \in (0, 1)$, with $\lambda_a + \lambda_b = 1$. Each agent is endowed with a cost of effort $c > 0$, independently drawn from a probability distribution that depends on group identity. We think of these “costs” as encompassing everything that hinders or helps individuals as they invest in skills: peer and neighborhood effects, innate ability, quality of schooling, resourcefulness of parents, and so on. (Thus, we consider it a

⁵Chan and Eyster (2003) have shown that lotteries can be used to pursue affirmative action goals when racial identity is not contractible. Fryer, Loury, and Yuret (2008) generalize this result and, using data on US college admissions, go on to estimate the efficiency losses from adopting blind rather than sighted preferential policies at the ex post stage. Fryer and Loury (2005b) study blind handicapping in the context of winner-take-all tournaments.

⁶A hybrid policy environment is also conceivable: one that is sighted ex ante but blind ex post. The view—popular in some circles in the United States—that using race in admissions to institutions of higher education is always wrong (blindness ex post) but that resources can legitimately be expended to narrow a racial performance gap in secondary schools (sightedness ex ante) illustrates this hybrid approach. For arguments consistent with this hybrid view, see the work of Thernstrom and Thernstrom (1997, 2003). One might imagine that there must also be a “sighted ex post, blind ex ante” hybrid scenario. But we show below that ex ante visibility is irrelevant to efficient policy design when the regulator is allowed to be sighted ex post.

reasonable assumption to take the population distribution of these costs as being predetermined from the point of view of the affirmative action policy maker.)

Should an agent acquire a slot, she will produce widgets at a rate that is equal to her stochastically determined individual productivity, whereas agents who do not purchase a slot are assumed to earn a payoff that is independent of productivity and that, with no further loss of generality, we normalize to zero.

Economic activity takes place in two stages in our model. In the ex ante stage, agents, distinguished by identity and effort cost (i, c) , choose whether or not to invest in skills, $s \in \{0, 1\}$. For example, this “investment” could be deciding to enter a training program or to crack the books in high school, where the costs can be thought of as an inverse measure of the agent’s endowed capacities. In the ex post stage, these same agents, distinguished now by identity and productivity (i, μ) , have a chance to purchase a slot at some price, p , after which production takes place. By “productivity” we simply mean the agents’ acquired ability to make use of their purchased slot. Clearly, an agent will purchase a slot at the ex post stage if and only if that agent’s productivity, μ , is no less than the market price, p . Thus, given our assumptions, we can express the utility of an agent making investment choice s , realizing productivity μ , and facing the slots price, p , as follows: $U = \max\{\mu - p, 0\} - cs$.

We take it that there is a fixed measure of slots, $\theta < 1$, and that slots are sold to the highest bidders in a competitive market.⁷ So, the agent whose productivity falls at the $(1 - \theta)$ th quantile of the population’s ex post productivity distribution will be the marginal buyer, and slots will trade at a price equal to this buyer’s valuation. The two stages of the model are linked because productivity is taken to be a noisy function of prior skill investment at the individual level, so a population’s distribution of productivity at the ex post stage is determined by the rate at which individuals in that population chose to make the skill-enhancing investment at the ex ante stage. Figure 1 illustrates the sequence of actions we envision.

B. Notation and Preliminaries

The primitives of our model are the distributions of agents’ costs and productivities. Let $G_i(c)$ be the probability that a group i agent is endowed with an effort cost that is less than or equal to c . The expression $G(c) \equiv \lambda_a G_a(c) + \lambda_b G_b(c)$ denotes the effort cost distribution for the entire population, with continuous density functions $g_i(c)$ and $g(c)$, respectively. These functions are assumed to be smooth and continuous with a common, connected support. The inverse functions, $G_i^{-1}(x)$ and $G^{-1}(x)$, $x \in [0, 1]$, give the effort cost of an agent at the x th quantile of the respective populations. By the law of large numbers and our continuum of agents assumption, $G_i(c)$ is also the fraction of agents in group i with effort cost less than or equal to c .

⁷The inelastic supply of slots assumption is innocuous and can be relaxed without affecting our results. That our slots are bought and sold on a competitive market is an abstraction, of course. In most settings in which affirmative action policies are employed, such as college admissions, the scarce positions for which agents compete are allocated administratively and are not auctioned to the highest bidders. However, our framing of the problem makes the economic intuition clear without sacrificing any meaningful generality. Furthermore, all our results can be shown to obtain under a planning/mechanism design formulation of this problem in which one assumes that agents’ ex ante stage investment costs and ex post stage productivity levels are private information.

We assume that group B is *disadvantaged* relative to group A , in the following sense.

Assumption 1. $g_a(c)/g_b(c)$ is a strictly decreasing function of c .

Monotonicity of this likelihood ratio implies that, for c interior to the cost support, (1) $G_a(c) > G_b(c)$, (2)

$$\frac{G_a(c)}{G_b(c)} > \frac{g_a(c)}{g_b(c)} > \frac{1 - G_a(c)}{1 - G_b(c)},$$

and (3) $G_a(c)/G_b(c)$ and $[1 - G_a(c)]/[1 - G_b(c)]$ are both strictly decreasing functions of c .

Given skill investment choice, $s \in \{0, 1\}$, let $F_s(\mu)$ be the probability that an agent's ex post productivity is no greater than μ . These distributions are also assumed to be smooth and continuous with a common, connected support. Their continuous density functions are denoted by $f_s(\mu)$. For a mass of agents with the common skill level s , $F_s(\mu)$ is the fraction of that mass with productivity less than or equal to μ . Investing in skills raises stochastic productivity in the following sense.

Assumption 2. $f_1(\mu)/f_0(\mu)$ is a strictly increasing function of μ . As before, monotonicity implies that, for μ interior to the productivity support, (1) $F_1(\mu) < F_0(\mu)$, (2)

$$\frac{F_1(\mu)}{F_0(\mu)} < \frac{f_1(\mu)}{f_0(\mu)} < \frac{1 - F_1(\mu)}{1 - F_0(\mu)},$$

and (3) $F_1(\mu)/F_0(\mu)$ and $[1 - F_1(\mu)]/[1 - F_0(\mu)]$ are both strictly increasing functions of μ .

Consider now the distribution of productivity in a population in which the fraction of agents who invest in skills is π . For π and μ , define

$$F(\pi, \mu) \equiv \pi F_1(\mu) + (1 - \pi) F_0(\mu).$$

Let $f(\pi, \mu)$ be the density and define the inverse, $F^{-1}(\pi, x)$, $x \in [0, 1]$, as the productivity level at the x th quantile of this ex post distribution.

Finally, let π_i be the fraction of group i agents who invest in skills, and let π denote the fraction of the overall population who invest. Then, $F(\pi_i, \mu)$ is the ex post distribution of productivity in that group, and

$$\lambda_a F(\pi_a, \mu) + \lambda_b F(\pi_b, \mu) = F(\lambda_a \pi_a + \lambda_b \pi_b, \mu) \equiv F(\pi, \mu)$$

is the corresponding productivity distribution for the population as a whole. Given assumption 2, it is easily verified that $\pi_a > \pi_b$ implies $f(\pi_a, \mu)/f(\pi_b, \mu)$ is a strictly increasing function of μ .

C. Market Equilibrium

Given the preliminaries above, the fraction of ex post stage agents with productivity greater than p can be written as

$$1 - F(\pi, p) = \pi[1 - F_1(p)] + (1 - \pi)[1 - F_0(p)].$$

This is the fraction of the population who demand a slot in the market, as a function of the price of slots, given π . Since the supply of slots is fixed, by assumption, at θ , it follows that the ex post stage equilibrium price, p^e —expressed as a function of the previously determined investment rate, π , and the inelastic slots supply, θ —satisfies $1 - F(\pi, p^e) = \theta$ or, equivalently, $p^e = F^{-1}(\pi, 1 - \theta)$.

Now, suppose that agents anticipate an ex post stage price of slots equal to p . Consider their ex ante stage investment decisions. Only the agents with cost c that is no greater than the anticipated net benefit from investment, which we denote $B(p)$, will acquire skills. At the ex post stage, an agent with productivity μ facing the slots price p will enjoy the net surplus of $\max\{\mu - p; 0\}$. Skills-enhancing investment alters ex ante stage productivity for an agent by causing it to be drawn at random according to the distribution $F_1(\mu)$ instead of the distribution $F_0(\mu)$. Thus (integrating by parts), we conclude that the net benefit of investment in skills can be written as follows:

$$B(p) \equiv \int_p^\infty (\mu - p)[f_1(\mu) - f_0(\mu)]d\mu = \int_p^\infty [F_0(\mu) - F_1(\mu)]d\mu \equiv \int_p^\infty \Delta F(\mu)d\mu,$$

where $F(\mu) - F_0(\mu)$ denotes the increment caused by investing in skills to the probability that an agent's ex post productivity exceeds μ . It follows that the fraction, π , of the ex ante stage population that acquires skills when the ex post stage slots price is anticipated to be p satisfies $\pi = G(B(p))$ or, equivalently, $G^{-1}(\pi) = \int_p^\infty \Delta F(\mu)d\mu$.

Combining these ex ante and ex post stage conditions, we conclude that (1) the laissez-faire equilibrium rate of skill acquisition, π^e , is determined by

$$G^{-1}(\pi^e) = \int_{F^{-1}(\pi^e, 1 - \theta)}^\infty \Delta F(\mu)d\mu \quad (1)$$

and (2) the laissez-faire equilibrium price of slots is given as

$$p^e = F^{-1}(\pi^e, 1 - \theta).$$

D. Social Efficiency

To characterize social efficiency under laissez-faire, define an *allocation* to be a pair of functions, $s(i, c) \in \{0, 1\}$ and $a(i, v) \in [0, 1]$, which specify the investment choice and the slot assignment probability for each type of agent, at the ex ante and ex post stages,

respectively. An allocation is *feasible* if it assigns a mass of agents to slots that does not exceed θ :

$$\sum_{i=a,b} \lambda_i \int_0^\infty a(i, \mu) f(\pi_i, \mu) d\mu \leq \theta, \quad (2)$$

where $\pi_i \equiv \int_0^\infty s(i, c) g_i(c) dc$ is the group i effort rate. And, it is *socially efficient* if it maximizes net social surplus:

$$\text{surplus} \equiv \sum_{i=a,b} \lambda_i \left[\int_0^\infty \mu a(i, \mu) f(\pi_i, \mu) d\mu - \int_0^\infty c s(i, c) g_i(c) dc \right], \quad (3)$$

among all feasible allocations.

To solve this maximization problem defining efficiency, let the fraction of agents exerting effort in some allocation be $\pi \in [0, 1]$. Efficiency requires that only the top θ productivity quantiles be assigned to slots and only the bottom π cost quantiles exert effort. So, the aggregate of widget values in this allocation can be no greater than $\int_{1-\theta}^1 F^{-1}(\pi, z) dz$ and the aggregate of effort costs can be no less than $\int_0^\pi G^{-1}(z) dz$. We conclude that this allocation can be efficient only if its effort rate equals π^* , where

$$\pi^* \equiv \arg \max \left\{ \int_{1-\theta}^1 F^{-1}(\pi, z) dz - \int_0^\pi G^{-1}(z) dz \right\}. \quad (4)$$

Below, we demonstrate that the laissez-faire market equilibrium is socially efficient. The intuition is straightforward: in any laissez-faire market equilibrium, the price, p^e , induces an agent to invest if and only if her cost is no greater than the benefit of investment, $B(p)$. Moreover, as the proof below makes clear, at the laissez-faire equilibrium slots price, the private benefit of skills-enhancing investment exactly equals its social marginal value.

Proposition 1. Under laissez-faire, $\pi^* = \pi^e$.

Proof. For an interior solution, the socially optimal rate of skill acquisition, π^* , implies the following necessary and sufficient first-order condition:

$$G^{-1}(\pi^*) = \int_{1-\theta}^1 \frac{\partial F^{-1}(\pi^*, z)}{\partial \pi} dz = \int_{F^{-1}(\pi^*, 1-\theta)}^\infty \Delta F(\mu) d\mu, \quad (5)$$

where we have used the implicit function theorem,

$$\frac{\partial F^{-1}(\pi, z)}{\partial \pi} = \left[\frac{\Delta F(\mu)}{f(\pi, \mu)} \right]_{\mu = F^{-1}(\pi, z)},$$

and the change of variables, $z \equiv F(\pi, \mu)$, $dz \equiv f(\pi, \mu) d\mu$, to derive the second equality.

From this it follows that the equilibrium skill investment rate under laissez-faire, π^e , which is determined by the equation

$$G^{-1}(\pi^e) = \int_{F^{-1}(\pi^e, 1-\theta)}^{\infty} \Delta F(\mu) d\mu,$$

and the socially optimal rate, π^* , are characterized by the same relationship. Note that the left-hand side in this equation, $G^{-1}(\pi)$, increases with π , while the right-hand side, $\int_{F^{-1}(\pi, 1-\theta)}^{\infty} \Delta F(\mu) d\mu$, decreases with π . So, the equation can have at most one solution, and thus $\pi^* = \pi^e$, which completes the proof.⁸ QED

Proposition 1 demonstrates that the laissez-faire equilibrium is socially efficient. The intuition is straightforward: given that there are no externalities to skill acquisition, when the slots are priced competitively, individual returns to skill acquisition are equal to social returns. Yet, owing to their predetermined social handicap, B 's will be less than proportionally represented among slot owners. This follows directly from the fact that under laissez-faire, A 's and B 's face the same return from acquiring skill, $B(p^e)$, but are characterized by different cost distributions, $G_i(c)$. Indeed, the group-specific skill acquisition and slot acquisition rates in equilibrium (π_i^e and σ_i^e , respectively) may be written as follows:

$$\pi_i^e = G_i(G^{-1}(\pi^e))$$

and

$$\sigma_i^e = 1 - F(\pi_i^e, F^{-1}(\pi^e, 1 - \theta)), i = a, b.$$

Therefore, the fact that group B is disadvantaged ($G_a(c) > G_b(c)$) implies that $\pi_b^e < \pi^e < \pi_a^e$, and therefore, $\sigma_a^e > \theta > \sigma_b^e$. That is, B 's invest in skills at the ex ante stage and acquire slots at the ex post stage at a lower rate than A 's in equilibrium. In the following section, the problem of designing an efficient policy to raise the representation of B 's is considered.

⁸An interior solution with $0 < \pi^* = \pi^e < 1$ obtains as long as the distribution of investment costs in the population spans a sufficiently wide range, in the sense that

$$G^{-1}(0) < \int_{F_0^{-1}(1-\theta)}^{\infty} \Delta F(\mu) d\mu$$

and

$$\int_{F_1^{-1}(1-\theta)}^{\infty} \Delta F(\mu) d\mu < G^{-1}(1),$$

which we implicitly assume.

III. Affirmative Action

Suppose one wants to decrease the group inequality that comes about in the laissez-faire equilibrium. How then should a policy maker intervene? And when? And with what effect?

In our model, it is possible to intervene at either the ex ante or the ex post stage and in either a “sighted” or “blind” manner. Sightedness implies the ability to tailor an intervention to group membership directly, for instance, different admissions requirements for Asians and whites. Blindness requires that the policy treat individuals without regard to their group while still achieving a target level of increased diversity, for example, selecting applicants on the basis of height while realizing that, on average, whites are taller than Asians. We treat each of these cases in turn.

A. Sighted Affirmative Action

Let σ_b denote a target rate of slot acquisition for group B agents, which is given. Assume that this target rate is greater than the above-specified laissez-faire rate for B 's, σ_b^e , but no greater than parity with A 's. Since the aggregate supply of slots is fixed at θ , a target slot acquisition rate for B 's of $\sigma_b \in (\sigma_b^e, \theta]$ implies a complementary target rate for A 's, $\sigma_a \in [\theta, \sigma_a^e)$, where

$$\sigma_a = \frac{\theta - \lambda_b \sigma_b}{\lambda_a}.$$

Namely, the rates at which the members of both groups are to acquire slots are predetermined by the capacity constraint and the representation target for the disadvantaged group.

It is trivial to show that a socially efficient sighted allocation that achieves these group-specific target rates of slot acquisition must have the property that, within each group, only the most productive agents will occupy the slots allocated to members of that group. Furthermore, only the lowest-cost agents within each group will acquire skills. (This reflects the key feature of sighted affirmative action: that it permits development and assignment policies to be formulated independently for each of the two groups.)

Reasoning as before—but now with respect to the two groups separately—we can characterize the socially efficient allocation of resources under sighted affirmative action in the following, straightforward manner: With σ_a and σ_b given, there exist group-specific skill investment rates, π_i , $i = a, b$, such that group i agents invest in skills if and only if their cost c does not exceed $G_i^{-1}(\hat{\pi}_i)$, and group i agents acquire slots if and only if their ex post productivity μ is no less than $F^{-1}(\pi_i, 1 - \sigma_i)$, where

$$\hat{\pi}_i \equiv \arg \max \left\{ \int_{1-\sigma_i}^1 F^{-1}(\pi, z) dz - \int_0^{\pi} G_i^{-1}(z) dz \right\}, i=a, b.$$

In light of our analysis of the laissez-faire case, we know that these rates $\hat{\pi}_i$ are uniquely defined by the following first-order conditions:

$$G_i^{-1}(\hat{\pi}_i) = \int_{1-\sigma_i}^1 \frac{\partial F^{-1}(\hat{\pi}_i, z)}{\partial \pi} dz = \int_{F^{-1}(\hat{\pi}_i, 1-\sigma_i)}^{\infty} \Delta F(\mu) d\mu, i=a, b.$$

Notice the similarities between the equations above and (5).

Now, it is readily verified that this constrained-efficient allocation can be implemented via the simple policy of providing a fixed, ex post subsidy to B 's for slot acquisition, as long as ex ante agents correctly anticipate that they will receive the full (subsidy-inclusive) rents associated with a competitive allocation of slots ex post. (Note that these rents will vary with both productivity and group identity.) Indeed, an ex post stage subsidy to B 's for slot acquisition in the amount $\hat{\tau} \equiv F^{-1}(\hat{\pi}_a, 1 - \sigma_a) - F^{-1}(\hat{\pi}_b, 1 - \sigma_b)$ ensures the desired representation of B 's among slot owners.

Given that the subsidy to B 's for slot acquisition is $\hat{\tau}$ and given that the supply of slots is fixed at θ , then it is readily seen that the ex post slots market clears at the subsidy-exclusive price of $p \doteq F^{-1}(\hat{\pi}_a, 1 - \sigma_a)$. If all agents correctly anticipate receiving their subsidy-inclusive rents ex post, then ex ante group-specific skill investment rates, $\pi_i, i = a, b$, would have to satisfy

$$\pi_a = G_a(B(\hat{p})) \text{ and } \pi_b = G_b(B(\hat{p} - \hat{\tau})),$$

assuming that there are no other market interventions. That is, $\pi_a = \hat{\pi}_a$ and $\pi_b = \pi_b$, implying that no explicit ex ante intervention influencing skill investment rates are needed to achieve social efficiency under sighted affirmative action.

The above derivation shows that when the ex post stage policy can be sighted, and given some target representation rate in the slots market for the disadvantaged group that is greater than their rate under laissez-faire but no greater than population proportionality, then no further intervention at the ex ante stage is required to achieve a constrained socially efficient allocation under affirmative action.

Furthermore, we have also shown that the optimal ex post sighted intervention can be implemented by subsidizing the acquisition of slots by members of the disadvantaged group at a fixed rate. This subsidy based on social identity raises the net price paid for slots in the ex post stage by group A agents and lowers the net price paid by group B agents. It thereby induces marginal members of the advantaged group to exit the slots market and marginal members of the disadvantaged group to enter it. This ex post subsidy also raises the ex ante skill acquisition rate among the disadvantaged while lowering it among the advantaged, relative to the laissez-faire outcome, in the unique socially efficient way.

To see this, recall that the following equations uniquely define the optimal group-specific skill investment rates, $\hat{\pi}_i$:

$$G_i^{-1}(\hat{\pi}_i) = \int_{1-\sigma_i}^1 \frac{\partial F^{-1}(\hat{\pi}_i, z)}{\partial \pi} dz.$$

Denote the solutions by the functions $\hat{\pi}_i(\sigma_i)$. Importantly, for any group-specific target, the optimal investment rate is completely determined by these functions. This, in turn, fully characterizes the ex post group-specific skill distribution if investment is optimal:

$$\hat{F}_i(\mu) \equiv F(\hat{\pi}_i(\sigma_i), \mu) = \hat{\pi}_i(\sigma_i) \times F_1(\mu) + [1 - \hat{\pi}_i(\sigma_i)] \times F_0(\mu).$$

The following observation is key: In any ex post equilibrium, the group-specific, market-clearing subsidy-inclusive prices, p_i , must satisfy $1 - F(\hat{\pi}_i(\sigma_i), p_i) = \sigma_i$, or $p_i(\sigma_i) = F^{-1}(\hat{\pi}_i(\sigma_i), 1 - \sigma_i)$. And, if agents in the groups anticipate these equilibrium prices, then investment at the ex ante stage is such that $G_i^{-1}(\hat{\pi}_i(\sigma_i)) = B(p_i(\sigma_i)) + t_i$, where t_i is a possible group-specific ex ante investment subsidy. It follows that if the allocation is constrained efficient, then the investment subsidy rates t_i must satisfy

$$t_i = G_i^{-1}(\hat{\pi}_i(\sigma_i)) - B(F^{-1}(\hat{\pi}_i(\sigma_i), 1 - \sigma_i)).$$

However, from the necessary conditions for social efficiency, we know that

$$G_i^{-1}(\hat{\pi}_i(\sigma_i)) = \int_{1-\sigma_i}^1 \frac{\partial F^{-1}(\hat{\pi}_i(\sigma_i), z)}{\partial \pi} dz = B(F^{-1}(\hat{\pi}_i(\sigma_i), 1 - \sigma_i)),$$

from which it follows that $t_i = 0$.

The economic intuition here is straightforward. Under a sighted ex post stage policy regime, the availability of slots to the members of each identity group is predetermined by the capacity and representation constraints. Thus, there are actually two separate allocation problems in the case of sighted affirmative action, one for each group. These separate problems are formally quite similar to the laissez-faire allocation problem, which we have already studied, with the difference being that the relevant availability of slots and distribution of agents' skill enhancement costs are specified separately for each of the two social identity groups. As a result, one can simply apply the analysis establishing the efficiency of laissez-faire equilibrium twice, separately for each of the two groups, in order to obtain our stated results: If when making their ex ante investment decisions members of each group correctly perceive the increment to ex post rents that results from skill enhancement and if these rents are the result of a competitive allocation of slots in the presence of the optimal ex post subsidy for the disadvantaged group, then one must neither tax nor subsidize skill acquisition to implement the constrained-efficient allocation since the expected return to skill enhancement by the members of both groups will be equal to its marginal social value.

B. Blind Affirmative Action

Our analysis of blind affirmative action proceeds in two steps. First, we take the ex ante group-specific skill investment rates, π_a and π_b , as given, presuming that $\pi_a > \pi_b$ and that, as before, the ex post stage target slot acquisition rates are fixed at σ_a and σ_b . We then solve for the constrained-optimal, blind, ex post stage slot allocation policy that maximizes gross output while achieving the target slot acquisition rates for each group. Second, using the indirect payoff function implied by the solution for this ex post constrained optimization problem, we solve for the group-specific skill investment rates, π_a and π_b , which maximize net social surplus while respecting the ex ante blindness constraint.

1. Optimal Ex Post Stage Policy—To begin, let some arbitrary group-specific ex ante skill investment rates $\pi_a > \pi_b$ be given and let the representation targets be fixed at σ_a and σ_b . Denote by $\pi \equiv \lambda_a \pi_a + \lambda_b \pi_b$ the population skill acquisition rate, and recall that $f(\pi, \mu)$ is the density of the ex post productivity distribution in a population of which the fraction π have invested ex ante.

A blind ex post stage policy is some function $\{a(\mu)\}_{\mu \geq 0}$ that gives the probability that an agent with productivity μ is assigned to a slot. Notice that this probability is independent of social identity. In the analysis that follows we restrict attention to *monotonic* policies, that is, functions $\{a(\mu)\}$ that are nondecreasing in μ .⁹ Define the function $\xi(\mu)$ as follows:

$$\xi(\mu) \equiv \frac{\lambda_b f(\pi_b, \mu)}{f(\pi, \mu)}.$$

The function $\xi(\mu)$ is the conditional probability that an agent observed to have ex post stage productivity μ actually belongs to group *B*. For $\pi_a > \pi_b$ the assumed monotone likelihood ratio property (MLRP; $f_1(\mu)/f_0(\mu)$ increasing with μ) implies that $\xi(\mu)$ is decreasing with μ ; that is, the more productive agents are less likely to belong to group *B*.

Given these definitions, the socially efficient blind ex post slot assignment policy, denoted $\{\bar{a}(\mu)\}$, is characterized by the following infinite-dimensional linear program:

$$\{\bar{a}(\mu)\}_{\mu \geq 0} \equiv \arg \max_{\{a(\mu)\}} \left\{ \int_0^\infty \mu a(\mu) f(\pi, \mu) d\mu \right\} \text{ subject to } \int_0^\infty \xi(\mu) a(\mu) f(\pi, \mu) d\mu = \lambda_b \sigma_b \in (\lambda_b \sigma_b^e, \lambda_b \theta] \quad (6)$$

(representation constraint),

$$\int_0^\infty a(\mu) f(\pi, \mu) d\mu = \theta$$

(capacity constraint), and

⁹If agents' productivities were treated as private information, as in a mechanism design formulation of this problem, then this monotonicity property would follow as a direct consequence of incentive compatibility.

$a(\mu) \in [0, 1]$ is nondecreasing in μ .

In what follows, we provide a solution for this optimization problem when $f_1(\mu)/f_0(\mu)$ is concave in μ . This ensures that $\xi(\mu)$ is convex.¹⁰

Before doing so, we digress to exhibit an example of productivity distributions $f_0(\mu)$ and $f_1(\mu)$ that generate a function $\xi(\mu) \equiv \lambda_b f(\pi_b, \mu)/f(\pi, \mu)$ that satisfies our requirement of being strictly decreasing and convex. Let productivity be distributed on the interval $[0, \bar{\mu}]$, and suppose that the density of the productivity distribution is $f_0(\mu) = a$ when there is no investment; if an agent invests in skills, then this density is $f_1(\mu) = h\mu$. (Since there are probability densities, we must have that $a = 1/\bar{\mu}$ and $h = 2/\bar{\mu}^2$.) These density functions are depicted in figure 2. Obviously, in this case the likelihood ratio

$$\frac{f_1(\mu)}{f_0(\mu)} = \frac{(2/\bar{\mu}^2)\mu}{1/\bar{\mu}} = \frac{2}{\bar{\mu}}\mu$$

and

$$\frac{\partial}{\partial \mu} \left[\frac{f_1(\mu)}{f_0(\mu)} \right] = \frac{2}{\bar{\mu}} > 0,$$

so our MLRP assumption is satisfied. Consider now the function $v(\mu) \equiv f(\pi_b, \mu)/f(\pi, \mu)$, which governs the shape of $\xi(\mu)$. Since

$$f(\pi, \mu) = \pi f_1(\mu) + (1 - \pi) f_0(\mu) = \pi \frac{2}{\bar{\mu}^2} \mu + (1 - \pi) \frac{1}{\bar{\mu}},$$

¹⁰To see this, define $r(\mu) \equiv f_1(\mu)/f_0(\mu)$ and rewrite $\xi(\mu)$ as

$$\xi(\mu) = \frac{\lambda_b f(\pi_b, \mu)}{f(\pi, \mu)} = \lambda_b \frac{\pi_b r(\mu) + (1 - \pi_b)}{\pi r(\mu) + (1 - \pi)}.$$

Therefore,

$$\xi'(\mu) = \lambda_b \frac{(\pi_b - \pi) r'(\mu)}{[\pi r(\mu) + (1 - \pi)]^2}$$

and

$$\xi''(\mu) = \lambda_b \frac{(\pi_b - \pi) \{ r''(\mu) [\pi r(\mu) + (1 - \pi)] - 2\pi [r'(\mu)]^2 \}}{[\pi r(\mu) + (1 - \pi)]^3}.$$

Since $\pi_b < \pi$, we have that $\xi''(\mu) > 0$ whenever $r''(\mu) < 0$.

we have that

$$v(\mu) \equiv \frac{\pi_b \frac{2}{\mu^2} \mu + (1 - \pi_b) \frac{1}{\mu}}{\pi \frac{2}{\mu^2} \mu + (1 - \pi) \frac{1}{\mu}}$$

and

$$\frac{\partial v(\mu)}{\partial \mu} = \frac{\frac{2}{\mu^3} (\pi_b - \pi)}{\left[\pi \frac{2}{\mu^2} \mu + (1 - \pi) \frac{1}{\mu} \right]^2} < 0$$

(recall that $\pi_b < \pi < \pi_a$). Moreover,

$$\frac{\partial^2 v(\mu)}{\partial \mu^2} = \frac{8}{\mu^5} \pi (\pi - \pi_b) \left[\pi \frac{2}{\mu^2} \mu + (1 - \pi) \frac{1}{\mu} \right]^{-3} > 0.$$

We conclude that

$$\frac{\partial \xi(\mu)}{\partial \mu} = \lambda_b \frac{\partial v(\mu)}{\partial \mu} < 0$$

and

$$\frac{\partial^2 \xi(\mu)}{\partial \mu^2} = \lambda_b \frac{\partial^2 v(\mu)}{\partial \mu^2} > 0,$$

so in this case $\xi(\mu)$ is indeed a convex function, as was to be shown.

We turn now to the task of solving the above-stated optimization problem, assuming that $\xi(\mu)$ is convex (which, as stated previously, is a condition on the ratio $f_1(\mu)/f_0(\mu)$). Toward this end, define $p(\bar{\pi}_a, \pi_b)$ as follows:

$$\frac{F(\pi_a, \bar{p})}{F(\pi_b, \bar{p})} \equiv \frac{1 - \sigma_a}{1 - \sigma_b}. \quad (7)$$

Under MLRP, with $\pi_a > \pi_b$, we know that $F(\pi_a, \mu)/F(\pi_b, \mu)$ is an increasing function of μ , so $\bar{p}$ exists and is uniquely defined by (7) whenever $\sigma_a > \sigma_b$.¹¹ Moreover,

$$\frac{F(\pi_a, \bar{p})}{F(\pi_b, \bar{p})} = \frac{1 - \sigma_a}{1 - \sigma_b} > \frac{1 - \sigma_a^e}{1 - \sigma_b^e} = \frac{F(\pi_a, p^e)}{F(\pi_b, p^e)}.$$

Thus, $\bar{p}$ is necessarily greater than the laissez-faire ex post stage slot market-clearing price, p^e :

$$\bar{p}(\pi_a, \pi_b) > p^e \equiv F^{-1}(\pi, 1 - \theta).$$

Using the implicit function theorem (and MLRP), one may verify that

$$\frac{\partial \bar{p}}{\partial \pi_a} = \frac{(1 - \sigma_b) \Delta F(\bar{p})}{(1 - \sigma_b) f(\pi_a, \bar{p}) - (1 - \sigma_a) f(\pi_b, \bar{p})} > 0$$

and

$$\frac{\partial \bar{p}}{\partial \pi_b} = \frac{-(1 - \sigma_a) \Delta F(\bar{p})}{(1 - \sigma_b) f(\pi_a, \bar{p}) - (1 - \sigma_a) f(\pi_b, \bar{p})} < 0.$$

Finally, given π and p , define the number $\bar{a}$ such that, if the measure of slots $\bar{a}$ were to be given away to the population at random while the remaining measure of slots $(\theta - \bar{a})$ were sold to the highest remaining bidders, then the market-clearing price for these remaining slots would be p . That is, let $\bar{a}$ be defined by the following equation:

$$\text{demand} = (1 - \bar{a})[1 - F(\pi, \bar{p})] \equiv \theta - \bar{a} = \text{supply}.$$

It follows that

$$\theta \geq \bar{a} \equiv 1 - \frac{1 - \theta}{F(\pi, \bar{p})} = 1 - \frac{F(\pi, p^e)}{F(\pi, \bar{p})} > 0.$$

Now, consider figure 3, which depicts the putatively optimal ex post blind assignment policy, denoted (μ) , along with an arbitrary alternative feasible policy, $a(\mu)$. The policy (μ) entails a random assignment of some free slots, for which all agents are eligible, together with a price floor (above the laissez-faire slots market-clearing price) at which losers of the lottery can purchase a slot with certainty. Specifically, all agents receive a free slot with a

¹¹To verify, define the function $l(p) = F(\pi_a, \pi)/F(\pi_b, p)$. Since

$$\frac{F(\pi_a, p^e)}{F(\pi_b, p^e)} = \frac{1 - \sigma_a^e}{1 - \sigma_b^e} < \frac{1 - \sigma_a}{1 - \sigma_b},$$

we have that $l(p^e) < (1 - \sigma_a)/(1 - \sigma_b)$. Furthermore, since $l(p) \rightarrow 1$ as $p \rightarrow \infty$, $(1 - \sigma_a)/(1 - \sigma_b) < 1$, and $l(p)$ is an increasing function of p , we have that there exists some $p > p^e$ such that $l(p) > (1 - \sigma_a)/(1 - \sigma_b)$. Therefore, the continuity of $l(p)$ guarantees a solution to eq. (7); i.e., there exists $p \in [p^e, p]$ such that

$$l(\bar{p}) = \frac{F(\pi_a, \bar{p})}{F(\pi_b, \bar{p})} = \frac{1 - \sigma_a}{1 - \sigma_b}.$$

The strict monotonicity of $l(p)$ further guarantees that this solution is unique.

probability $\in (0, \theta)$, while agents with $\mu > p$ who do not win the lottery successfully bid for one of the remaining slots in the open market and pay the price p .

More formally, we have the following result.

Theorem 1. Given the group-specific ex ante investment rates, $\pi_a > \pi_b$, under the MLRP assumptions on $F_e(\mu)$ and $G_t(c)$, and with the additional assumption that $f_1(\mu)/f_0(\mu)$ is concave in μ , then, for $\bar{\mu}$ and $\bar{p}$ as defined above, the solution to the ex post linear optimization problem under blind affirmative action, $\{\bar{a}(\mu)\}_{\mu \geq 0}$, is given as follows:

$$\bar{a}(\mu) = \begin{cases} \bar{a} \equiv 1 - \frac{1-\theta}{F(\pi, \bar{p})} & \text{if } \mu \leq \bar{p} \\ 1 & \text{if } \mu > \bar{p}. \end{cases}$$

Proof. Define $\bar{a}(\mu)$ as above and let $a(\mu)$ be any alternative feasible assignment policy. We begin by using first principles to restrict the set of alternative policies that we need to consider. The putatively optimal policy $\bar{a}(\mu)$ assigns the full measure θ of slots, so any alternative policy $a(\mu)$, with $a(\mu) \leq \bar{a}(\mu)$ everywhere and with strict inequality somewhere in the support of the distribution of μ , must violate the capacity constraint. Also, any alternative policy $a(\mu)$ that lies everywhere below $\bar{a}(\mu)$ could never be optimal. For such a policy would leave some slots unassigned and could be perturbed in a way that respects the representation constraint while using the unassigned slots to increase widget production. Finally, notice that any alternative policy $a(\mu)$ with the property that $\min_{\mu} (a(\mu) - \bar{a}(\mu)) > 0$, with strict inequalities somewhere in the support of the distribution of μ , could never be optimal. As figure 3 makes clear, an alternative policy of this sort, when compared to $\bar{a}(\mu)$, would necessarily shift assignment weight from more productive agents to less productive agents, thereby lowering aggregate widget production. Combining these observations, we conclude that, without any loss of generality, we can restrict attention to monotonic alternative policies $a(\mu)$ that cross $\bar{a}(\mu)$ from below. (Such an alternative policy is depicted in fig. 3.) Let μ' denote the first point of crossing. Obviously, $a(\mu) < \bar{a}(\mu)$ for $\mu < \mu'$, $a(\mu) = \bar{a}(\mu)$ for $\mu \in (\mu', p]$, and $a(\mu) > \bar{a}(\mu)$ for $\mu \in (p, \mu_{\max}]$.

Next, we assert that there exists a point, denoted $\hat{\mu} \in (\mu', p)$, such that

$$\int_0^{\hat{\mu}} [a(\mu) - \bar{a}(\mu)] f(\pi, \mu) d\mu = 0 = \int_{\hat{\mu}}^{\mu_{\max}} [a(\mu) - \bar{a}(\mu)] f(\pi, \mu) d\mu.$$

To see this, notice that the continuous function

$$H(\mu) \equiv \int_0^{\mu} [a(x) - \bar{a}(x)] f(\pi, x) dx$$

is decreasing for $\mu < \mu'$ and for $\mu > p$, while it satisfies

$$H(0) = \lim_{\mu \uparrow \mu_{\max}} H(\mu) = 0.$$

So a point $\hat{\mu} \in (\mu', p)^-$ must exist at which

$$H(\hat{\mu}) = \int_0^{\hat{\mu}} [a(x) - \bar{a}(x)] f(\pi, x) dx = 0.$$

Now, define the functions $\phi(\mu)$ and $\psi(\mu)$ as follows:

$$\phi(\mu) \equiv [a(\mu) - \bar{a}(\mu)] f(\pi, \mu)$$

and

$$\psi(\mu) \equiv \xi(\hat{\mu}) + [\mu - \hat{\mu}] \xi'(\hat{\mu}) - \xi(\mu).$$

Since $\xi(\mu)$ is convex, $\psi(\mu)$ is a nonpositive concave function that achieves its maximum of zero at $\mu = \hat{\mu}$ (see fig. 4). Moreover, feasibility of the two policies and the foregoing arguments imply

$$\int_0^\infty \xi(\mu) \phi(\mu) d\mu = 0 = \int_0^{\hat{\mu}} \phi(\mu) d\mu = \int_{\hat{\mu}}^\infty \phi(\mu) d\mu.$$

Therefore, consulting figure 5 and after some manipulation, one can see that

$$\begin{aligned} & \int_0^\infty \phi(\mu) \psi(\mu) d\mu \\ &= \xi'(\hat{\mu}) \int_0^\infty \mu [a(\mu) - \bar{a}(\mu)] dF(\pi, \mu) \equiv \xi'(\hat{\mu}) \Delta Q \\ &= \int_0^{\hat{\mu}} \phi(\mu) \psi(\mu) d\mu \\ &+ \int_{\hat{\mu}}^\infty \phi(\mu) \psi(\mu) d\mu \geq \psi(\mu') \int_0^{\hat{\mu}} \phi(\mu) d\mu \\ &+ \psi(\bar{p}) \int_{\hat{\mu}}^\infty \phi(\mu) d\mu = 0, \end{aligned}$$

where Q denotes the change in gross output occasioned by moving to the policy $a(\mu)$ from the putatively optimal policy $\bar{a}(\mu)$. The inequality in the third line uses the properties of the function $\psi(\mu)$ that are stated above and the fact that $a(\mu) - \bar{a}(\mu) < 0$ for $\mu \in [0, \mu']$, $a(\mu) - \bar{a}(\mu) = 0$ for $\mu \in [\mu', p]$, and $a(\mu) - \bar{a}(\mu) > 0$ for $\mu \in [p, \infty)$. Now, since $\xi'(\hat{\mu}) < 0$, it follows that $Q \leq 0$. QED

To get some intuition for the result stated in theorem 1, notice, as the diagram in figure 3 illustrates, that any deviation from the putatively optimal policy $\{\bar{a}(\mu)\}$ shifts assignment weight to the middle of the range of realized productivity values from the extremes. Feasibility requires that this shift must preserve the rate at which group B agents are assigned to slots (e.g., it must preserve the conditional expected value of $\xi(\mu)$ among those assigned to slots). So it is intuitively clear (see fig. 4) that when $\xi(\mu)$ is a convex function,

less weight needs to be shifted into the middle of the range of μ from below than is shifted into the middle of that range from above, in order to preserve the conditional expected value of $\xi(\mu)$ among those assigned to slots when moving from $\{(\mu)\}$ to $\{a(\mu)\}$. Thus, the expectation of μ under the alternative policy must fall after these shifts.¹²

2. Ex Ante Stage Policy

With this result in hand, we can now characterize the optimal ex ante policy under blind affirmative action. We will show that this optimal development policy entails a universal subsidy for (respectively, tax on) skill investment whenever B 's are more prevalent on the development (respectively, assignment) margin. Using the theorem above and our assumptions, we can write the ex post indirect payoff function, denoted $Q(\pi_a, \pi_b; \sigma_a, \sigma_b)$, as follows ($Q(\cdot)$ is the maximal aggregate widget output ex post under blind affirmative action that is obtainable from the available slots, given that ex ante group-specific skill investment rates are π_a and π_b and that ex post assignment target rates are σ_a and σ_b). With $\bar{\mu}(\pi) \equiv \int_0^\infty \mu f(\pi, \mu) d\mu$ denoting the mean productivity in the ex post stage population when the fraction π have invested in skills in the ex ante stage and letting

$$\mu^+(\pi, p) \equiv \frac{\int_p^\infty \mu f(\pi, \mu) d\mu}{1 - F(\pi, p)}$$

denote the conditional mean productivity among agents who are no less productive than p , we can see from the theorem that constrained optimal gross widget production ex post under blind affirmative action is given by (suppressing dependence of these functions on π_a, π_b, σ_a , and σ_b)

$$\bar{Q} = [\bar{a}] \bar{\mu}(\pi) + [\theta - \bar{a}] \mu^+(\pi, \bar{p}) = [\bar{a}] \bar{\mu}(\pi) + [1 - \bar{a}] \int_{\bar{p}}^\infty \mu f(\pi, \mu) d\mu,$$

where $\bar{a} = 1 - [(1 - \theta)/F(\pi, p)]$ and where $\bar{p} \equiv \bar{p}(\pi_a, \pi_b; \sigma_a, \sigma_b)$ is such that

$$\frac{F(\pi_a, \bar{p})}{F(\pi_b, \bar{p})} \equiv \frac{1 - \sigma_a}{1 - \sigma_b}.$$

We conclude that the optimal ex ante policy is one that chooses skill investment rates for the groups π_a and π_b , in a manner consistent with the ex ante blindness constraint, so as to maximize $Q(\pi_a, \pi_b; \sigma_a, \sigma_b)$.

Let us now consider this ex ante maximization problem. If a policy is blind at both stages, all agents face the same return from investing in skills ex ante, no matter what subsidies or taxes may be levied on skill acquisition or slot purchases. Hence, if the overall skill

¹²Optimal blind policy takes on a qualitatively similar form even without the assumptions on the shape of $f_1(\mu)/f_0(\mu)$. In particular, the problem can be reformulated so that it becomes (the dual of) what Anderson and Nash (1987, sec. 4.4) call a continuous semi-infinite linear program. Applying their theorem 4.8 (p. 76)—which makes explicit use of Caratheodory's theorem—one concludes that optimal policy under blind affirmative action can be expressed as a step function. See Chan and Eyster (2003, proof of proposition 2) for a similar argument.

investment rate in the population were π and if only the lowest-cost agents are investing (as is required for optimality), then the cost of the marginal investor in either group would have to be $G^{-1}(\pi)$. It follows that, under blindness, the group-specific rates of ex ante skill investment rates, π_a and π_b , must be related to the aggregate investment rate π in the following way: $\pi_i = G_i(G^{-1}(\pi))$, $i = a, b$. Hence, the socially optimal ex ante stage aggregate skill investment rate under blind affirmative action, denoted $\bar{\pi}$, is given by

$$\bar{\pi} \equiv \arg \max_{\pi \in [0,1]} \left\{ [\bar{a}] \bar{\mu}(\pi) + [1 - \bar{a}] \int_{\bar{p}}^{\infty} \mu f(\pi, \mu) d\mu - \int_0^{\pi} g^{-1}(z) dz \right\},$$

where $\bar{a} \equiv 1 - [(1 - \theta)/F(\pi, p)]$, $\bar{p} \equiv p(\bar{\pi}_a, \bar{\pi}_b; \sigma_a, \sigma_b)$, and $\pi_a = G_a(G^{-1}(\pi))$; $\pi_b = G_b(G^{-1}(\pi))$.

Computing the derivatives, we have the following expressions for $d\bar{\pi}/d\pi$ and $d\bar{p}/d\pi$:

$$\frac{d\bar{a}}{d\pi} = \left[\frac{1 - \bar{a}}{F(\pi, \bar{p})} \right] \left[f(\pi, \bar{p}) \frac{d\bar{p}}{d\pi} - \Delta F(\bar{p}) \right], \quad \frac{d\bar{p}}{d\pi} = \Delta F(\bar{p}) \left[\frac{(1 - \sigma_b) \frac{g_a(c)}{g(c)} - (1 - \sigma_a) \frac{g_b(c)}{g(c)}}{(1 - \sigma_b) f(\pi_a, \bar{p}) - (1 - \sigma_a) f(\pi_b, \bar{p})} \right],$$

for $c \equiv G^{-1}(\pi)$ —the ex ante cost of the marginal skill investor (the “development margin”)—and for $p \equiv p(G_a(c), G_b(c); \sigma_a, \sigma_b)$ —the ex post market-clearing slots price (the “assignment margin”).

The first-order necessary condition for an aggregate investment rate π to solve this maximization problem may be written as follows:

$$\begin{aligned} G^{-1}(\bar{\pi}) &= \bar{a} \frac{d\bar{\mu}}{d\pi} + (1 - \bar{a}) \left\{ \int_{\bar{p}}^{\infty} \mu [f_1(\mu) - f_0(\mu)] d\mu - \bar{p} f(\pi, \bar{p}) \frac{d\bar{p}}{d\pi} \right\} \\ &+ \left(\frac{d\bar{a}}{d\pi} \right) \int_0^{\bar{p}} \mu f(\pi, \mu) d\mu \\ &= \bar{a} \int_0^{\infty} \mu [f_1(\mu) - f_0(\mu)] d\mu \\ &+ (1 - \bar{a}) \int_{\bar{p}}^{\infty} [\mu \\ &- \bar{p}] [f_1(\mu) \\ &- f_0(\mu)] d\mu \\ &+ \left[\frac{1 - \bar{a}}{f(\pi, \bar{p})} \right] [\Delta f(\bar{p}) \\ &- f(\pi, \bar{p}) \frac{d\bar{p}}{d\pi}] \int_0^{\bar{p}} (\bar{p} \\ &- \mu) f(\pi, \mu) d\mu. \end{aligned}$$

Now, to see whether implementing this optimal ex ante investment rate requires a tax or a subsidy on skill acquisition, consider the incentive for an agent to invest in skill ex ante,

given that this agent anticipates the ex post policy (π, p) . (To be clear, we are referring here to the policy “receive a slot for free with probability π and the opportunity to purchase a slot at price p with probability $1 - \pi$.”) If this agent fully and correctly anticipates the rents associated with the allocation of positions under this policy at the ex post stage, then her benefit from skill investment, $B(\pi, p)$, is given by

$$B(\pi, p) = \pi \int_0^\infty \mu [f_1(\mu) - f_0(\mu)] d\mu + (1 - \pi) \int_p^\infty [\mu - p] [f_1(\mu) - f_0(\mu)] d\mu.$$

So, in the absence of any subsidy for or tax on ex ante skill acquisition, the fraction π of the population investing when anticipating the ex post policy (π, p) would be $\pi = G(B(\pi, p))$. By comparing $B(\pi, p)$ to the expression for $G^{-1}(\pi)$ derived in the first-order condition above, we see that implementing the constrained optimal blind ex ante investment rate, π , requires a universal subsidy for (a tax on) skill acquisition whenever $F(p) > (<) f(\pi, p) \frac{d\pi}{dp}$, that is, whenever

$$1 > (<) \frac{(1 - \sigma_b) \frac{g_a(\bar{c})}{g(\bar{c})} - (1 - \sigma_a) \frac{g_b(\bar{c})}{g(\bar{c})}}{(1 - \sigma_b) \frac{f(\pi_a, \bar{p})}{f(\pi, \bar{p})} - (1 - \sigma_a) \frac{f(\pi_b, \bar{p})}{f(\pi, \bar{p})}},$$

where $c \equiv G^{-1}(\pi)$. Using the facts that

$$\begin{aligned} f(\pi, \mu) &\equiv \lambda_a f(\pi_a, \mu) + \lambda_b f(\pi_b, \mu), \\ g(c) &\equiv \lambda_a g_a(c) + \lambda_b g_b(c), \\ \lambda_a \sigma_a + \lambda_b \sigma_b &= \theta \end{aligned}$$

and that the denominator above is positive (which follows from the MLRP), we conclude that the optimal ex ante stage policy requires a subsidy for (tax on) skill investments whenever

$$\frac{g_b(\bar{c})}{g(\bar{c})} > (<) \frac{f(\pi_b, \bar{p})}{f(\pi, \bar{p})}.$$

That is, under blind affirmative action the optimal ex ante policy entails a universal skill subsidy whenever B 's are relatively more prevalent on the development margin than they are on the slot assignment margin.

IV. Concluding Remarks

Affirmative action policy involves interventions designed to promote the presence in scarce positions of the members of disadvantaged social identity groups. In this paper, we write down a simple economic model to explore the efficient design of such policies. Our key insights derive from the interrelationship that naturally exists between the “early” and the “late” interventions that may be undertaken to reduce inequality between identity groups. Our notion of “early” pertains to the development of persons through the acquisition of

productive traits, while “late” alludes to the assignment of persons whose productivities are predetermined to scarce positions that give them the opportunity to produce. We write down a tractable model for both of these stages and examine how the optimal intervention to reduce group inequality plays out across the development and the assignment margins.

In our simple and transparent laboratory—two jobs, supply-demand allocation late, rational skill acquisition early, and a heterogeneous population— we establish three basic results. All of these results turn on the fact that an inescapable behavioral connection exists between the development and assignment stages: on the basis of how positions are allocated, net aggregate output and position-specific rents arise *ex post*. The anticipation of these rents induces agents to incur a cost in order to acquire productive traits. Thus, the anticipation of *ex post* rents drives outcomes *ex ante*.

Our first result is that the *laissez-faire* outcome is socially efficient but produces unequal group representation in desired positions. With competitive slots pricing *ex post*, the incremental rent a person expects from *ex ante* skill enhancement is precisely equal to its marginal social value.

Our second result is that when identity is fully visible and contractible, the efficient affirmative action policy avoids explicit human capital subsidies for the disadvantaged. This seemingly counterintuitive result follows immediately once the problem has been posed in our simple framework. To prefer a group of people at the assignment stage of the production process is already to implicitly subsidize their acquisition of skills. If these implicit benefits are correctly anticipated by the agents and if the *ex post* intervention is itself efficient, then no further interference with investment incentives is desirable.

This, we show, is no longer the case when preferential policies must be identity-blind. This is our third result. Even with an efficient *ex post* policy, private and social returns from *ex ante* investment will generally not coincide under blindness. We have shown that at the *ex post* stage, the second-best, identity-blind intervention to increase opportunity for a disadvantaged group involves allocating some productive opportunities for free and at random while allowing the remaining slots to be acquired in the open market at a price above the *laissez-faire* level. We have also derived an empirically meaningful and ultimately testable condition under which the efficient blind *ex ante* policy entails a general subsidy to human capital investments.

Acknowledgments

We are grateful to Pierre-Andre Chiappori, Edward Glaeser, Lawrence Katz, Debraj Ray, Bernard Salanie, and two anonymous referees for helpful comments and suggestions. We also thank seminar participants at numerous universities. Loury gratefully acknowledges financial support from the Andrew W. Mellon Foundation and the Carnegie Corporation.

References

- Anderson, Edward J.; Nash, Peter. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. New York: Wiley; 1987.
- Bowen, William; Bok, Derek. *The Shape of the River*. Princeton, NJ: Princeton Univ. Press; 1998.

- Chan, Jimmy; Eyster, Erik. Does Banning Affirmative Action Lower College Student Quality? *A.E.R.* 2003; 93(2):858–872.
- Coate, Stephen; Loury, Glenn C. Will Affirmative Action Policies Eliminate Negative Stereotypes? *A.E.R.* 1993; 83(5):1220–1240.
- Deshpande, Ashwini. Affirmative Action in India and the United States. World Development Report. Washington, DC: World Bank; 2006. 2006
- Fryer, Roland; Loury, Glenn C. Affirmative Action and Its Mythology. *J. Econ. Perspectives.* 2005a; 19(3):147–162.
- Fryer, Roland; Loury, Glenn C. Affirmative Action in Winner-Take-All Markets. *J. Econ. Inequality.* 2005b; 3(3):263–280.
- Fryer, Roland; Loury, Glenn C.; Yuret, Tolga. An Economic Analysis of Color-Blind Affirmative Action. *J. Law, Econ., and Org.* 2008; 24(2):319–355.
- Galanter, Mark. Competing Equalities: Law and the Backward Classes in India. Oxford: Oxford Univ. Press; 1992.
- Holzer, Harry; Neumark, David. Assessing Affirmative Action. *J. Econ. Literature.* 2000; 38(3):483–568.
- Klitgaard, Robert. Elitism and Meritocracy in Developing Countries. Baltimore: Johns Hopkins Univ. Press; 1986.
- Loury, Glenn C. The Anatomy of Racial Inequality. Cambridge, MA: Harvard Univ. Press; 2002.
- Lundberg, Shelly; Startz, Richard. Private Discrimination and Social Intervention in Competitive Labor Markets. *A.E.R.* 1983; 73(3):340–347.
- Moro, Andrea; Norman, Peter. Affirmative Action in a Competitive Economy. *J. Public Econ.* 2003; 87(3–4):567–594.
- Sowell, Thomas. Affirmative Action around the World: An Empirical Study. New Haven, CT: Yale Univ. Press; 2004.
- Thernstrom, Abigail; Thernstrom, Stephen. America in Black and White: One Nation, Indivisible. New York: Simon & Schuster; 1997.
- Thernstrom, Abigail; Thernstrom, Stephen. No Excuses: Closing the Racial Gap in Learning. New York: Simon & Schuster; 2003.
- Welch, Finis. Employment Quotas for Minorities. *J.P.E.* 1976 Aug; 84(4, pt. 2):S105–S139.

Fig. 1.
Sequence of actions

Fig. 2.
An example of productivity distributions that generate $\xi(\mu)$ convex

Fig. 3. Uniqueness of optimal stage 2 policy: any alternative feasible policy $\alpha(\mu)$ must shift weight to the middle from the extremes of the range of μ , relative to (μ) .

Fig. 4.
Uniqueness of equilibrium under a convex likelihood function

Fig. 5.