MIDS WORKING PAPER NO. 234

Meritocracy in the Face of Group Inequality

Rajiv Sethi
Barnard College, Columbia University, New York

Rohini Somanathan
Delhi School of Economics, University of Delhi, Delhi

June 2019

Madras Institute of Development Studies

79, Second Main Road, Gandhi Nagar, Adyar, Chennai 600 020
Tel.: +91 44 2441 2589   pub@mids.ac.in   www.mids.ac.in

Meritocracy in the Face of Group Inequality*

Rajiv Sethi† Rohini Somanathan‡

Abstract

Meritocratic systems are commonly understood as those that assign tasks to individuals who can best perform them. But future performance cannot be known prior to assignment, and must be inferred from other traits. We consider a model in which performance depends on two attributes—ability and training—where ability is endowed and unobserved and training is acquired and observed. The potential to acquire training depends on ability and resource access, so ability affects performance through two channels: indirectly through training and directly through the performance function. The population consists of two identity groups, each with the same ability distribution, but with differential access to resources. We characterize the sets of training levels that maximize expected performance. An allocation is *monotonic* if, for each group, there is a threshold value of training such that all those above this value (and none below) are selected. It is *group-blind* if assignment is independent of group identity, and *psuedomeritocratic* if it is both monotonic and group-blind. We show that performance-maximizing allocations are not generally monotonic or group-blind, and are pseudomeritocratic under only very special conditions. This is true even when individuals can respond to non-monotonic policies by underinvesting in training, or when commitment to selection policies is possible.

---

* Delivered by Rohini Somanathan as a Public Lecture at the Madras Institute of Development Studies on 25 January 2018.

Acknowledgements: We thank Kehinde Ajayi, Nageeb Ali, V. Bhaskar, Michael Lachmann, Glenn Loury, François Maniquet, Debraj Ray, Abhiroop Sarkar, and S. Subramanian for helpful suggestions.

†Department of Economics, Barnard College, Columbia University, New York and the Santa Fe Institute, Santa Fe.
‡Department of Economics, Delhi School of Economics, University of Delhi, Delhi.

1 Introduction

Any practical implementation of a meritocratic ideal must be based on observables. In popular discourse, merit-based allocations are inconsistent with selection based on markers of group identity, such as ethnicity, gender, or religion, since these are not intrinsically related to performance. It is also believed that selection based on merit ought to satisfy a monotonicity property; those with higher values of performance-related attributes, such as a test scores or course grades, should have precedence over those with lower values. In this paper we challenge the general nature of both of these claims.

We are interested in the problem of allocating a set of prized or elite positions to a fraction of the population, and define a meritocratic assignment as one which assigns tasks to the individuals who would best perform them. Since future performance cannot be known at the time of selection, it has to be inferred from signals. As a consequence, attributes that are not intrinsically related to merit may nevertheless be informative, and be valid criteria for selection under a meritocratic policy. Furthermore, under quite general conditions, the inference made from a signal of merit need not have the monotonicity property: lower values may sometimes signal greater merit—in the sense of higher expected performance—than higher values.

We develop these arguments on the basis of a simple model with two groups. Within each group, individuals differ along a dimension we call *ability*, under which we include all those characteristics having the same distribution across groups. They also differ in their access to *resources* such as good schools, neighborhoods or conducive family environments. There are two levels of resource access, high and low, and one group is disadvantaged in the sense that that it has a smaller proportion of people with high resource access. Neither ability nor resources are independently observable, but they jointly determine an individual's potential to acquire *training*, which is observable, and which can therefore be used as a criterion for selection. Future performance depends on both ability and training, so ability operates through two channels: it affects performance directly, and also indirectly through its effect on training. A meritocratic allocation assigns the available positions to those with the highest expected performance.

We say that an allocation of individuals to positions is *monotonic* if, for each group, there is a threshold value of training such that all those above this value (and none below) are selected. It is *group-blind* if the likelihood of selection for an individual is not contingent on their group membership. And it is *psuedomeritocratic* if it is both monotonic and group-blind. Many college affirmative action programs have been challenged in court because they violate psuedomeritocracy in applying lower admission thresholds to minority candidates. We show that such violations may be consistent with allocations based on our notion of merit.

More precisely, we show that under quite general conditions, performance-maximizing allocations can fail to satisfy monotonicity. This happens when training is heavily resource dependent, so that some levels of training are difficult to attain for high ability individuals with low resource access. Since performance depends on both ability and training, maximization of expected performance may require the selection of individuals with lower training (and higher inferred ability) compared with some who have higher training.

Non-monotonic policies create incentives for some individuals to underinvest in training, so that they may increase their prospects for selection. As a result, a given level of observed training may involve the pooling of individuals with varying levels of potential training. However, given differences in resource access across groups, such pools will have different levels of inferred ability, and the performance-maximizing selection policy will not generally be group blind. In fact, the disadvantaged group can be favored under the optimal policy to such an extent that its members end up being overrepresented in elite positions. If commitment to a selection policy is possible, underinvestment can be avoided, but selection will still not generally be group blind, and will typically favor the disadvantaged group. This is the case even though neither diversity nor social justice are explicit criteria for selection.

These arguments reveal that the usual framing of the problem of affirmative action—as a trade-off between performance and representativeness—is misleading. Imposition of a pseudomeritocracy constraint can lower performance, and requiring more equal representation can have lower efficiency costs than requiring selection to be pseudomeritocratic.

There is a large literature on identity-contingent admission and hiring policies when there is imperfect information on the performance-related characteristics of potential candidates. Early work focused on statistical discrimination (Phelps, 1972; Arrow, 1973; Coate and Loury, 1993; Aigner and Cain, 1977; Cornell and Welch, 1996). More recently, there has been a focus on optimal selection rules when diversity and merit are both valued (Chan and Eyster, 2003; Fryer et al., 2008; Fryer and Loury, 2013). These papers ask how traditional or "sighted" affirmative action compares with color-blind affirmative action, which refers to policies that are not explicitly group-contingent, but are nevertheless motivated by diversity goals. In this literature, individual merit is treated as synonymous with some observable qualification such as a test score. As a result, a purely meritocratic allocation simply involves the application of a common qualification threshold to all members of the population, regardless of identity.[1]

The contributions of Durlauf (2008), Roemer (2009), and Scanlon (2018) are directly related to the questions we ask here and use similar conceptions of merit. In evaluating the efficiency implications of alternative assignment rules, Durlauf argues for a "move from merit as reward to merit as effectiveness" on the grounds that true meritocracy should be understood as being

[1]For instance, Chan and Eyster (2003, p.860) assume that the "expected academic promise... of a candidate with test score t is simply t: the higher a candidate's score, the higher her quality."

based on future potential rather than past achievement. He points out that this "means that merit needs to be assessed relative to the properties of human capital production functions." Along similar lines, Roemer argues that a proper conception of merit must be based on an individual's current attributes rather than on what has been achieved in the past, and Scanlon observes that merit depends heavily on context and institutional goals. This notion of meritocracy as expected performance-maximization is the one we adopt here.

Cestau et al. (2017) express the same idea in the context of the selection into a program for gifted schoolchildren (emphasis added):

> A student who has, in some way, experienced hardship may underperform on achievement tests relative to his or her capability. By taking account of such empirically grounded differences across demographic groups, a district may be better able to determine which students are most suited to admission to the gifted program... While this profiling based on differences in distributions across racial groups is beneficial to minority students, *it is not preferential treatment*.

Thus Cestau et al. (2017) distinguish between *profiling*, which is an attempt to use demographic information to better meet performance goals unrelated to diversity, and *affirmative action*, which involves preferential treatment for a group beyond levels justified by profiling. The former serves the goal of maximized performance, while the latter increases diversity relative to performance-maximizing levels. The authors consider selection decisions made by an (unnamed) school district, and find that the district engages in both profiling and affirmative action with respect to family income, and engages in profiling but not affirmative action with respect to race. One of our main points is that bans on the use of group membership in the process of selection—such as California's Proposition 209 or Michigan's Proposal 2—are bans on both affirmative action and performance profiling. Contrary to the rhetoric surrounding such initiatives, such mandates therefore block the implementation of meritocratic allocations.

There are other reasonable conceptions of meritocracy based on the distinction between factors we cannot control and those we can. For instance, Loury (1981) distinguishes between "weak meritocracy" (defined as a positive correlation between income and ability) and "strong meritocracy" (which requires that the conditional distribution of ability rises with income in the sense of first-order stochastic dominance). The long-run income distribution in his model satisfies weak meritocracy but can fail to satisfy strong meritocracy. The reason for this failure has some similarity to the mechanism at work in our analysis: there may be a range of income levels such that those near the top of the range have high training but low ability on average, while those below them have lower training but higher ability in expectation. But in Loury's model incomes are directly related to productivity, so the outcome is indeed meritocratic in the performance-related sense used here.

In the next section we present our model. In section 3, we consider a discrete version in which future performance can be fully inferred from training provided that all individuals attain the highest training levels consistent with their ability and resources. We identify conditions under which selection policies fail to be monotonic, resulting in underinvestment and pooling in equilibrium. Section 4 considers the case of continuous ability, which involves ability uncertainty even in the absence of underinvestment. We end with some reflections on how our results contribute to the discourse on efficiency, meritocracy and equal opportunity.

2 The Model

There is a continuum population composed of two groups, 1 and 2, with population shares s_1 and s_2 respectively. Individuals within each group differ in their ability a, and access to resources r. There is a common distribution of ability $F(a)$, but the distribution of resources varies by group. There are two resource levels, r_l and r_h where $r_l < r_h$. The proportion of individuals in group i with access to the higher resource level is denoted by q_i. We assume, without loss of generality, that $q_1 \leq q_2$ and refer to the first group as *disadvantaged* if the inequality is strict.

Ability and resources are independently distributed and neither can be directly observed. We observe only group identity and a signal t, which we refer to as *training*. This could be thought of as an observable measure of educational attainment or an indicator of it such as a test score. The highest attainable level of training for individual is given by the continuous function $\tau(a, r)$, which is increasing in both arguments. That is, at each level of ability and resources, the chosen level of training satisfies

$$t \leq \tau(a, r).$$

We are interested in the problem of selecting a fraction k of the population into scarce positions, and refer to k as *elite capacity*. Performance in these positions is increasing in both ability and training as given by the function $p = \phi(a, t)$. We therefore allow ability to have both a direct and an indirect effect on performance, capturing the idea that past scores might under-represent the capacity of talented individuals to perform.

Any selection policy can be described by a pair of functions $\pi_i(t)$ that denote the probability of being selected conditional on exhibiting training t and belonging to group $i \in \{1, 2\}$. If the policy is deterministic, with each individual being accepted or rejected with certainty, it can equivalently be described in terms of the group-contingent sets of training levels that result in selection.

Our goal is to characterize the selection policy that would assign the highest performers to the available positions. We can think of this problem one faced by a social planner or a single firm or university facing a fixed pool of applicants. Since the planner only observes training and group identity, he uses these to infer performance. The nature of these inferences depends on the form of the training and performance functions, $\tau(a, r)$ and $\phi(a, t)$, as well as the strategic choices of training levels by candidates in equilibrium.

We refer to a selection policy as *monotonic* if, for each group i, there is some threshold level of training such that selection is ensured for those with training above the threshold, and rejection is ensured for those below it. That is, under a monotonic selection policy, $\pi_i(t) > 0$ implies $\pi_i(t') = 1$ for all $t' > t$. A policy is *group-blind* if selection policies are independent of group membership: $\pi_1(t) = \pi_2(t)$ at all t.

If a policy is both monotonic and group-blind, we call it *pseudomeritocratic*. Such policies correspond to common notions of meritocracy as reward. Under a monotonic policy, no individual can increase her likelihood of selection by underinvesting in training. However, as we show below, performance-maximizing policies need not be monotonic, which creates incentives for underinvestment.

We define a meritocratic assignment as one which maximizes expected performance and study the nature of selection in such an assignment under three scenarios. In the first, individuals choose the highest training that they can achieve, given their ability and resources, independently of the selection policy. That is, we have $t = \tau(a, r)$ at all levels of ability and resource access. This would occur, for instance, if there was high consumption value to reaching one's potential. We refer to this as the baseline case.

Next, we allow for individuals to underinvest in training. The selection policy and the training distribution are now jointly determined in equilibrium, with individual investments in training being chosen to maximize the likelihood of being selected. That is, given an anticipated selection policy, individuals choose a training level that maximizes the probability of selection. If $t_i(a, r)$ denotes the training level chosen by someone in group i who has ability a and resource access r, we have

$$t_i(a, r) \in \arg\max_{t \leq \tau(a, r)} \pi_i(t).$$

An equilibrium in this case is a selection policy and a set of training choices such that expected performance is maximized given the training choices, and no individual can increase the likelihood of selection by adopting a different, feasible level of training.

Finally, we allow for commitment to a selection policy before training investments are made. This results in selection probabilities that are weakly increasing in training, so that there are no incentives for underinvestment. Expected performance is higher with commitment than without, but even with commitment polices will not generally be pseudomeritocratic.

3 Meritocracy as Effectiveness

We begin with a discrete version of the model in which individuals in each group have either high or low ability (a_l and a_h) and consider first the baseline case without underinvestment. That is, suppose that all individuals achieve their maximum training level, $\tau(a, r)$.

Since we have two resource levels, there are now four ability-resource combinations. The set of maximum achievable training levels for these is denoted by $T = \{t_{ll}, t_{hl}, t_{lh}, t_{hh}\}$, where t_{ij} for ability i and resource j. We know that those with low ability and low resources have the lowest level of both training and performance and those with high ability and high resources have the highest of both. The remaining two cases cannot be unambiguously ranked in terms of either training or performance.

If higher levels of training always signal higher ability, we say that the training function satisfies *inference monotonicity*. That is, inference monotonicity holds if $t_{lh} < t_{hl}$. Whether or not inference monotonicity holds will depend on the relative importance of ability and resources in determining training; if training is heavily resource dependent and resource differences are large, inference monotonicity can fail.

We say that *performance monotonicity* holds if higher levels of training signal higher expected performance. Clearly inference monotonicity is sufficient but not necessary for performance monotonicity. If $t_{hl} < t_{lh}$ but $\phi(a_h, t_{hl}) < \phi(a_l, t_{lh})$, then inference monotonicity fails but performance monotonicity nevertheless holds, and higher training levels will correspond to higher levels of performance.

Whether or not we have performance monotonicity depends on the relative importance of ability and training in determining performance; if training is sufficiently important we could have performance monotonicity even if inference monotonicity fails.

If elite capacity is sufficiently small, all positions can be filled by individuals with the highest training level. In this case there is no tension between meritocracy as reward and meritocracy as effectiveness. For larger levels of elite capacity this is no longer the case.

Suppose that no individual underinvests in training, performance monotonicity fails to hold, and elite capacity is greater than the share of the total population with the highest training level. Then performance-maximization requires that the selection policy itself be non-monotonic, skipping over the second highest training level to draw from the third. To illustrate, consider the following example.

Example 1. *Suppose that $t_{lh} > t_{hl}$ and $\phi(a_h, t_{hl}) > \phi(a_l, t_{lh})$, so performance monotonicity fails to hold. Model parameters are $s_1 = s_2 = 0.5, F(a_l) = 0.7, q_1 = 0.1, q_2 = 0.3$, and $k = 0.1$. Then if all individuals attain their highest feasible training level, the performance-maximizing selection policy will set $\pi_i(t_{hh}) = 1$*

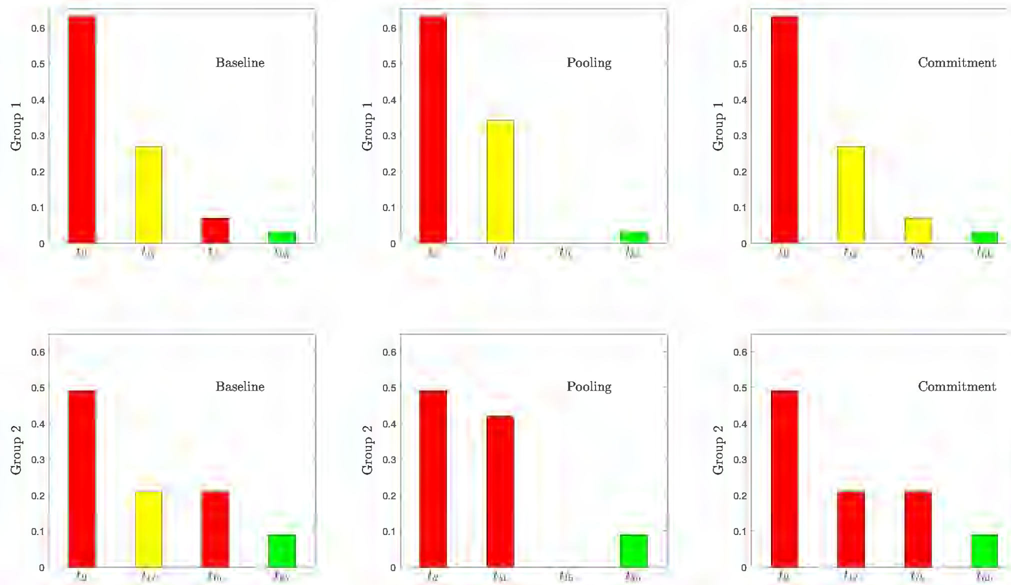

Figure 1: Optimal selection policies under three conditions.

and $\pi_i(t_{ll}) = \pi_i(t_{lh}) = 0$ *for both groups.*

In this example, of the 10% percent selected to elite positions, 6% will have training t_{hh}, while the remaining 4% will be drawn from those having training t_{hl}. This draw could (but need not) be uniform at random without regard to group membership, in which case the policy is blind to identity, but not monotonic in observable signals of performance.

This is illustrated in the left column of Figure 1, which shows the training frequency distributions in the two groups; green bars indicate certain selection, red indicate certain rejection, and yellow indicate selection with a probability that is positive but less than one. If those with training level t_{hl} are selected uniformly across the two groups, the disadvantaged group will be underrepresented in elite positions.

A non-monotonic selection policy of this kind clearly creates incentives for underinvestment in training. Unless the consumption benefits of attaining one's potential are sufficiently great, those with low ability and high resources will attempt to pool with those immediately below them, in order to increase the likelihood of being selected.

This underinvestment, however, will change the quality of the pool of candidates with training level t_{hl}, for two reasons: first, because individuals who were initially at training t_{lh} had lower performance than those with whom they are pooling, and second, because they are lowering their performance even further by underinvesting in training relative to their potential.

In order for complete pooling at t_{hl} to be consistent with equilibrium, it must be the case that the expected performance of the pool be higher than that of the individuals who were initially at training t_{lh}. The exact conditions for this are identified below, but for the purposes of the example, suppose that these conditions are satisfied in both groups. Then, under the conditions of Example 1, there is an equilibrium in which all individuals with low ability and high resources choose training level t_{hl} rather than t_{lh}, thus pooling with those at a training level below their potential. Despite being diluted, this pool has higher expected performance (in both groups) than those known to be of low ability and high resources. But since the pool in the disadvantaged group has higher expected performance, and is large enough to fill the elite positions remaining after those with the highest training level t_{hh} have been selected, we have $\pi_1(t_{hl}) > 0 = \pi_2(t_{hl})$ in equilibrium.

This is illustrated in the middle column of Figure 1. Specifically, in this example, half of those in the advantaged group pool at t_{hl} have high ability, while the corresponding figure for the disadvantaged group is more than three-quarters. Given the value of k, elite positions will be filled by selecting all those at the highest training level (regardless of group), together with some of those in the disadvantaged group pooled at t_{hl}.

Note that in this example the disadvantaged group will end up being *overrepresented* in elite positions, occupying 55% of elite positions, despite having a worse distribution of training, and an identical distribution of ability. And this effect arises even though the only criterion for selection is the maximization of expected performance.

If commitment to a selection policy were possible, a higher level of performance could be attained by simply applying the same probability of selection to those with training t_{hl} and t_{lh} even in the second case. This is shown in the right column of the Figure 1, for the same numerical example. In this case the optimal selection policy does not induce underinvestment. Note, however, that the composition of the selected pool will be the same as in the case without commitment, with advantaged group members selected if and only if they attain the highest level of training.

The example shown in Figure 1 reveals a robust phenomenon. To identify conditions under which this arises, define λ_i as follows:

$$\lambda_i = \frac{(1-q_i)(1-F(a_l))}{(1-q_i)(1-F(a_l)) + q_i F(a_l)}$$

This is the probability that an individual picked randomly from among those with training in the set $\{t_{hl}, t_{lh}\}$ is of high ability.

Note that if $q_1 < q_2$, then we must have $\lambda_1 > \lambda_2$. Among the pool of individuals with the two intermediate training levels, those in the disadvantaged group must have higher ability (though not necessarily higher performance) in expectation.

Next suppose that $t_{hl} < t_{lh}$ (so inference monotonicity fails to hold), and define $\mu \in (0,1)$ by the following condition:

$$\mu\phi(a_h, t_{hl}) + (1-\mu)\phi(a_l, t_{hl}) = \phi(a_l, t_{lh}).$$

Here μ is the proportion of high ability individuals in a pool with training t_{hl} that would result in the same expected performance as a pool of low ability individuals with training t_{lh}.

If $\lambda_2 > \mu$ then, provided that performance monotonicity fails, there is an equilibrium in which all individuals with low ability and high resources choose training level t_{hl} rather than t_{lh}, thus pooling with those at a training level below their potential. Despite being diluted, this pool has higher expected performance (in both groups) than those known to be of low ability and high resources. As long as the pool in the disadvantaged group has higher expected performance, and is large enough to fill the elite positions remaining after those with the highest training level t_{hh} have been selected, we have $\pi_1(t_{hl}) > 0 = \pi_2(t_{hl})$ in equilibrium.

As a step towards providing a more complete characterization of equilibrium policies, define k_1 as follows:

$$k_1 = (1 - F(a_l))(s_1 q_1 + s_2 q_2).$$

This is the proportion of the entire population having both high ability and high resources, and hence the largest level of elite capacity that can be filled with those with the highest level of training. Define k_2 as follows:

$$k_2 = 1 - F(a_l)(s_1(1-q_1) + s_2(1-q_2)).$$

This is the proportion of the total population with either high ability or high resources (or both).

It is clear that if if inference (and hence performance) monotonicity holds, then there exists an equilibrium with a pseudomeritocratic selection policy at all values of elite capacity. If performance monotonicity fails to hold, however, pseudomeritocratic selection policies can arise in equilibrium only if elite capacity is sufficiently small or sufficiently large (see the appendix for all proofs):

Proposition 1. *If performance monotonicity fails to hold, there is an equilibrium with a pseudomeritocratic selection policy if and only if $k < k_1$ or $k > k_2$.*

This result rules out equilibria with pseudomeritocratic allocations for intermediate levels of elite capacity. In this range, equilibrium selection policies have a more complicated structure, and involve underinvestment:

Proposition 2. *Suppose that performance monotonicity fails to hold, and $k \in (k_1, k_2)$. Then*

a) *If $\lambda_1 > \mu$, there is an equilibrium selection policy with $\pi_i(t_{ll}) = \pi_i(t_{lh}) = 0$, $\pi_i(t_{hh}) = 1$, and $\pi_1(t_{hl}) > \pi_2(t_{hl})$. There is underinvestment by all individuals with low ability and high resource access in the disadvantaged group, and at least some with low ability and high resource access in the advantaged group.*

b) *If $\lambda_1 < \mu$, there is a group-blind equilibrium selection policy with $\pi_i(t_{ll}) = 0$, $\pi_i(t_{hh}) = 1$, and $\pi_i(t_{hl}) = \pi_i(t_{lh}) < 1$. There is underinvestment by some individuals with low ability and high resource access in both groups.*

In the first case, the probability of selection must be strictly greater for those in the disadvantaged group at some training levels. All low ability high resource types in the disadvantaged group underinvest, pooling with those having high ability and low resource access. In the advantaged group such pooling may be partial if $\lambda_2 < \mu$, in which case low ability high resource types distribute themselves across the two intermediate training levels in such a manner as to equalize performance at μ.

In the second case, individuals in *both* groups distribute themselves across the two intermediate training levels in such a manner as to make expected performance equal to μ. This allows the equilibrium selection policy to be blind to group identity, though it violates monotonicity by applying the same probability (below one) to multiple training levels. As a result, there is underinvestment in equilibrium.

Proposition 2 implies that the equilibrium selection policy will favor the disadvantaged group for some values of elite capacity, provided that $\lambda_1 > \mu$. As we have see in the middle column of Figure 1, this effect can be strong enough to result in overrepresentation of the disadvantaged group. Figure 2 shows how representation varies with elite capacity for the same numerical example, under the assumption that $\lambda_1 > \mu$. We see here that overrepresentation of the disadvantaged group arises for a wide range of elite capacities, and such overrepresentation can be substantial.

While the case of two ability levels reveals quite clearly that a performance-maximizing selection policy can be both group-contingent and non-monotonic, it has an important limitation: the training level fully reveals ability in the absence of underinvestment. As a result, when all individuals choose their highest attainable training level, a group blind policy will be optimal. This is no longer the case when ability is distributed continuously, however, since there will be ability uncertainty even without underinvestment. We consider this next.

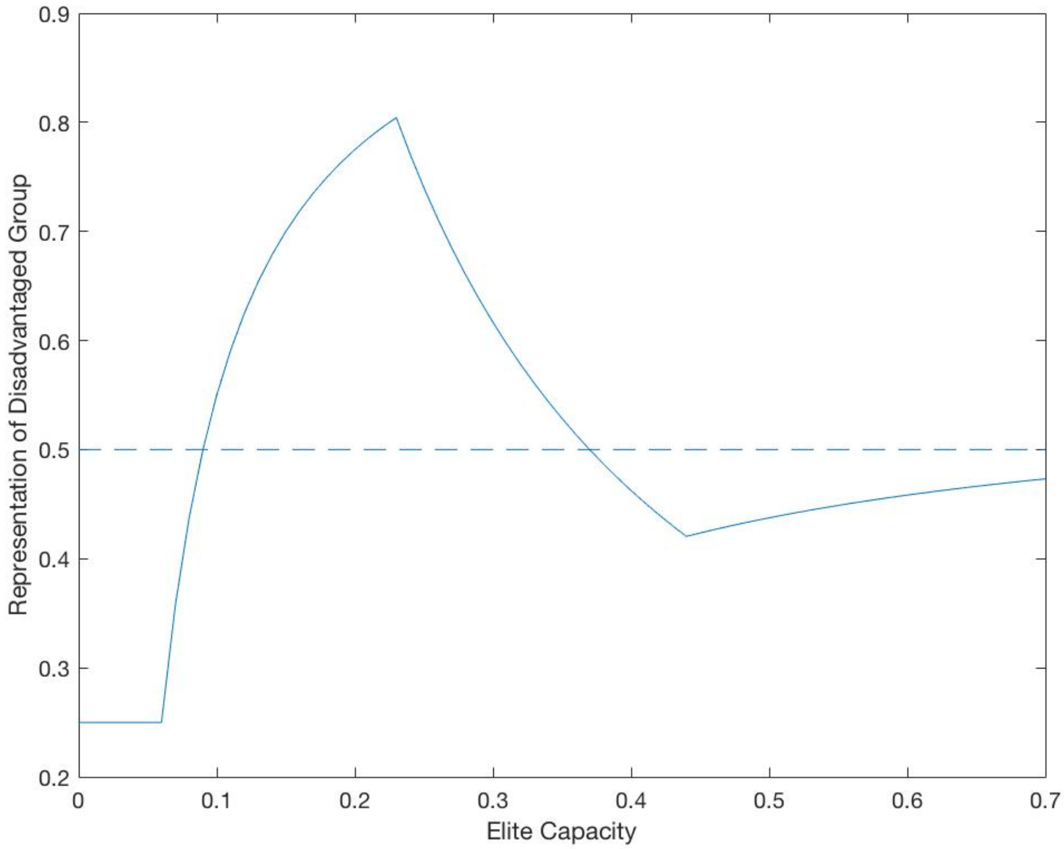

Figure 2: Disadvantaged Group Representation as Elite Capacity Varies

4 Continuous Ability Distributions

Now suppose that the distribution of ability is continuous with density $f(a)$ and support $[0,1]$. We assume that f is continuous and strictly positive at all points in the support.[2] As before, the ability distributions are the same in the two groups.

Let $\bar{t} = \tau(1, r_h)$ denote the highest attainable level of training, and $t^* = \tau(1, r_l)$ the highest level attainable by those with low resource access. Since τ is increasing in both ability and resource access, $t^* < \bar{t}$. As a result, there is a range of training levels which reveal an individual to have high resource access.[3]

As before, we first consider the baseline case without underinvestment in training, so $t =$

[2]The assumption that the density function is everywhere positive ensures that a likelihood function is well-behaved at all relevant points, and simplifies our analysis.

[3]This assumption greatly simplifies the argument for the baseline case with no underinvestment, since ability is also revealed when training is above t^*. The case $t^* = \bar{t}$ gives rise to qualitatively similar results, as discussed in the appendix.

$\tau(a, r)$. In this case the optimal selection policy will be generically deterministic, with each individual selected with probability zero or one.[4] To characterize the performance-maximizing selection policy, let T_1 and T_2 denote sets of group-contingent training levels, with the interpretation that individuals in group i are selected if and only if they have a training level in T_i. We restrict attention to selection sets that are chosen from the set $\mathcal{T}$ of finite unions of subintervals of $[0, \bar{t}]$.

Given any level of training t, let $\alpha_l(t)$ and $\alpha_h(t)$ denote the ability levels defined implicitly by

$$t = \tau(\alpha_l, r_l)$$

and

$$t = \tau(\alpha_h, r_h).$$

That is, $\alpha_l(t)$ is the ability level needed to attain training t if one has low resource access, and $\alpha_h(t)$ is the corresponding ability level for those with high resource access. Clearly $\alpha_l(t) > \alpha_h(t)$; any given level of training requires more ability to attain if one's resource access is low. Furthermore, both α_h and α_l are strictly increasing in t, since a higher threshold at any given level of resource access requires higher ability to meet.

Given any set of training levels $T \in \mathcal{T}$, let $A_h(T)$ and $A_l(T)$ denote the corresponding sets of ability levels for high and low resource access individuals. That is, $A_h(T)$ is the set of ability levels that result in training within the set T, provided that one has high resource access, and $A_l(T)$ is analogously defined. Specifically,

$$T = \tau(A_l, r_l)$$

and

$$T = \tau(A_h, r_h).$$

Note that since $T \in \mathcal{T}$ and τ is continuous, $A_h(T)$ and $A_l(T)$ will both be elements of $\mathcal{A}$, the set of finite unions of subintervals of $[0, 1]$.

Firms cannot observe resource access, but they can observe group membership, and since groups may have different levels of resource access, firms may choose T_1 and T_2 to be different. Given any such choice, define $m_h(T_i)$ and $m_l(T_i)$ as the mass of individuals in group i having ability levels in $A_h(T_i)$ and $A_l(T_i)$ respectively:

$$m_h(T_i) = \int_{A_h(T_i)} dF(a), \qquad m_l(T_i) = \int_{A_l(T_i)} dF(a)$$

Then expected performance is given by

$$E(p) = \sum_{i=1}^{2} s_i \left(\frac{q_i}{m_h(A_i)} \int_{A_h(T_i)} \phi(a, \tau(a, r_h)) dF + \frac{1 - q_i}{m_l(A_i)} \int_{A_l(T_i)} \phi(a, \tau(a, r_l)) dF \right). \tag{1}$$

[4]The reason why selection is not deterministic in Chan and Eyster (2003) despite an exogenously given score distribution is that the selection is assumed to monotonic. Relaxing this assumption results in deterministic selection, as in Ray and Sethi (2010). When we allow for under-investment in training, the performance-maximizing selection policy will no longer be deterministic in general.

The sets T_i are chosen to maximize this subject to the capacity constraint

$$\sum_{i=1}^{2} s_i \left(q_i m_h(T_i) + (1-q_i) m_l(T_i)\right) = k, \tag{2}$$

where $k \in (0,1)$ is the proportion of total jobs that are in the skilled sector.

We shall refer to a set of training levels (T_1, T_2) that satisfies the capacity constraint (2) as an *allocation*. An allocation that maximizes expected productivity (1) is a *performance-maximizing allocation*.

With continuous ability, an allocation is monotonic if

$$t \in T_i \implies t' \in T_i$$

for each i and each t, t' such that $t' > t$. That is, a monotonic performance-maximizing allocation has the property that each set T_i can be identified with a threshold training level t_i such that an individual in group i secures an elite position if and only if her training exceeds t_i.

Group-blind allocations satisfy $T_1 = T_2$, and allocations that are both monotonic and group-blind are pseudomeritocratic. In this case there is a common threshold $t_1 = t_2$ such that an individual secures an elite position if and only if her training exceeds this threshold, regardless of group membership. At any such allocation, the only relevant characteristic for judging merit is the value of the observable signal t, and not what this signal might imply, in conjunction with group identity, about the expected productivity of an applicant.

An allocation is *group egalitarian* is it satisfies the following representation target:

$$q_i \int_{A_h(T_i)} dF(a) + (1-q_i) \int_{A_l(T_i)} dF(a) = k, \tag{3}$$

for each i. This ensures that the proportion of individuals who secure elite positions is the same in both groups (and equal to the share of the total population in this sector). A group egalitarian allocation need not be monotonic, and indeed we show below that a performance-maximizing allocation can be group egalitarian and non-monotonic.

Much of our analysis is conducted under the assumption that $t^* < \bar{t}$, which follows from the hypothesis that τ is increasing in r when $a > 0$. This, together with the assumption that f is everywhere positive, implies that expected productivity conditional on training is falls discontinuously at t^*. Given this, non-monotonicity of the selection rule is ensured for one or both groups for certain values of elite capacity constraint. But this discontinuity is not essential to the argument, and we show by means of an extended example in the appendix that non-monotonicity can arise even if $t^* = \bar{t}$, in which case firms can never deduce resource access at any level of observed training.

4.1 Performance-Maximizing Allocations

A performance-maximizing allocation need not be monotonic because productivity depends directly on ability, as well as indirectly through its effect on training. Since resource access is heterogeneous, it is possible that at some training levels the population is composed largely of high ability and low resource individuals, while at higher training levels the opposite is true. Under these conditions, expected productivity may be higher at lower training levels. We now identify conditions under which this can occur.

Let $\gamma_i(t)$ denote the likelihood that an individual in group i with training t has high resource access. For $t > t^*$ we clearly have $\gamma_i(t) = 1$, since such training levels are unattainable for those with low resource access. For $t \leq t^*$ we have

$$\gamma_i(t) = \frac{q_i f(\alpha_h(t)) a_h'(t)}{q_i f(\alpha_h(t)) a_h'(t) + (1 - q_i) f(\alpha_l(t)) a_l'(t)}$$

The expected productivity of someone in group i with training t is then

$$E(p_i|t) = \gamma_i(t)\phi(\alpha_h(t), t) + (1 - \gamma_i(t))\phi(\alpha_l(t), t).$$

For training levels above t^*, firms face no uncertainty about applicant resource access. In this case expected productivity is simply

$$E(p_i|t) = \phi(\alpha_h(t), t).$$

Since both ϕ and α_h are common to both groups and increasing in t, it follows that for $t > t^*$, $E(p_1|t) = E(p_2|t)$, and $E(p_i|t)$ is increasing in t.

Expected productivity is not increasing in t everywhere, however. Since $\gamma_i(t^*) < 1$ for each i, there is uncertainty about the ability of an individual with training t^* regardless of group identity. As a result, individuals with maximal ability and low resource access are pooled with those having moderate ability and high resource access. Since ability matters for productivity, expected productivity at t^* exceeds that at levels of training slightly above t^*. This implies that performance-maximizing allocations must be non-monotonic for some levels of elite capacity, as we show below.

Let G_i denote the distribution function for training in group i, and set

$$\hat{k} = \sum_{i=1}^{2} s_i(1 - G_i(\hat{t}_1)) > 0. \tag{4}$$

We show below that if elite capacity lies below this threshold, then the performance-maximizing allocation must be monotonic, and indeed pseudomeritocratic.

To identify conditions under which the performance-maximization allocation is not monotonic, for each group i, let $\tilde{t}_i$ denote the *highest* level of training in the interval $(0, t^*)$ that satisfies the following condition:

$$E(p_i|\tilde{t}_i) = \phi(\alpha_h(t^*), t^*).$$

This exists because

$$E(p_i|0) < \phi(\alpha_h(t^*), t^*) < E(p_i|t^*).$$

Now define $\tilde{k}$ as the elite capacity level that would exactly absorb the population with training above these respective thresholds:

$$\tilde{k} = \sum_{i=1}^{2} s_i(1 - G_i(\tilde{t}_i)).$$

The following result describes how the performance-maximizing allocation varies with elite capacity in the baseline case.

Proposition 3. *Suppose no individual chooses to underinvest in training. Then the performance-maximizing allocation is pseudomeritocratic if $k < \hat{k}$, and non-monotonic if $k \in (\hat{k}, \tilde{k})$. Furthermore, there exists $k' \in [\tilde{k}, 1)$ such that, if $k > k'$, then the performance-maximizing allocation is monotonic.*

If elite capacity is sufficiently small, only those with high resource access can reach training levels that allow them to be selected. And among these, only those with the highest ability are selected. Since only elites are selected, and among these only those with highest ability, the same standard is applied to both groups.[5] If elite capacity is sufficiently large, we again obtain monotonicity of performance-maximizing allocations under weak conditions. This is obviously true if $k = 1$, and also true for k sufficiently large.

So if elite capacity is very small or very large, we obtain monotonicity. But non-monotonic performance-maximizing allocations arise under very weak conditions when elite capacity lies in a range that is neither too large not too small. The following example with two identical groups illustrates.

Example 2. *Suppose the two groups are of identical size and composition, $r_l = 1$, $r_h = 2$, $t = \sqrt{ar}$, and $p = \sqrt{at}$. Then $t^* = 1$ and $\bar{t} = \sqrt{2}$. The performance-maximizing allocation is not monotonic for $k \in (0.11, 0.51)$. When $k = 0.3$, the set of training levels that result in selection (in either group) is $T_i = [0.91, t^*] \cup [1.14, \bar{t}]$.*

Example 2, illustrated in Figure 3, shows that the performance-maximizing allocation may fail to be meritocratic in an intuitive sense even when it is group egalitarian. Among those with greater resource access, there are some with lower ability (and hence lower productivity) that get elite positions while others with higher ability and productivity are denied. This is because the former are pooled with high ability resource poor individuals, while the latter are not.

[5] Note, however, that if there are few elites in the disadvantaged group, then $\hat{k}$ will be small, and hence the requirement for such allocations will be restrictive.

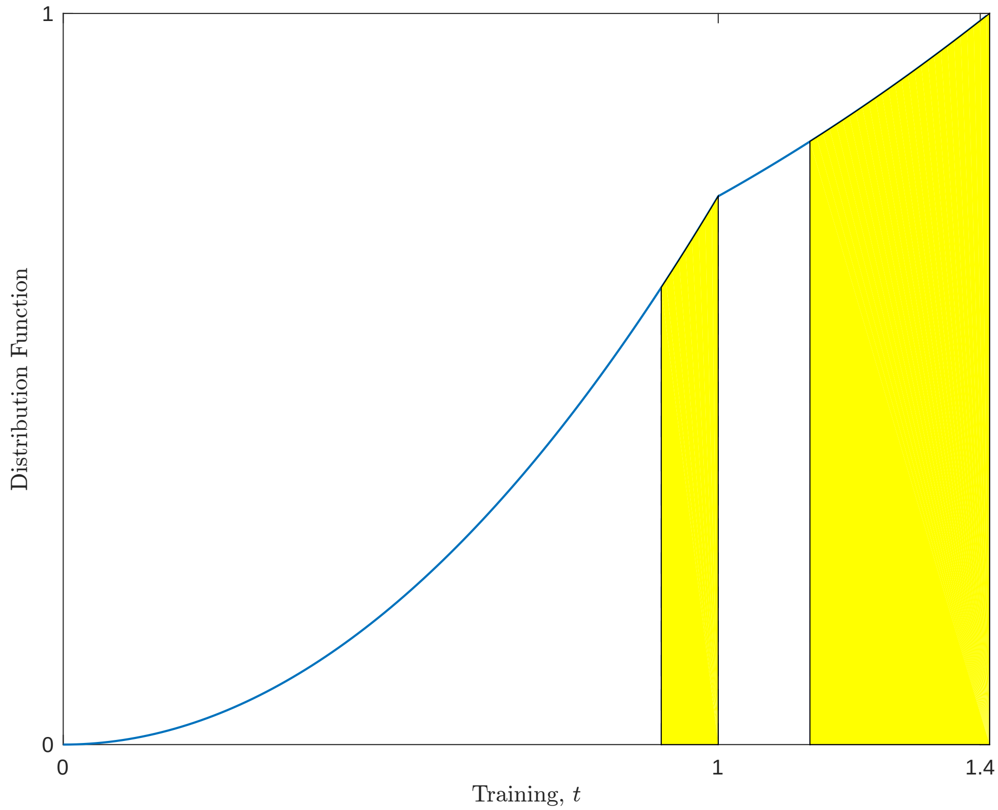

Figure 3: A Nonmonotonic Performance-maximizing Allocation with Identical Groups.

Since the two groups are identical in this example, the performance-maximizing allocation is group egalitarian. For groups that are not identical but very close in advantage, constraining firms to select a group egalitarian outcome will not entail much productivity loss. In contrast, requiring them to choose a pesudomeritocratic outcome can entail first order productivity losses if the performance-maximizing allocation is nonmonotonic. That is, under certain conditions, imposing group egalitarianism on firms will be less burdensome than imposing pseudomeritocracy.

4.2 Differential Treatment

We have seen that if elite capacity is small enough, the performance-maximizing allocation is pseudomeritocratic, with a common selection threshold for both groups. A disadvantaged group in this case will be underrepresented, so the outcome will not be group-egalitarian. Furthermore, *within* each group those with lower resource access will be underrepresented.

For intermediate values of elite capacity the performance-maximizing allocation is not-monotonic and therefore also not pseudometritocratic. But this raises the possibility that belonging to a disadvantaged group might make selection *more* likely at some levels of training. The reason is that group membership is observable while ability is not, and belonging to a group with lower resource access on average signals higher ability at any given level of training.

This is most easily seen if a is uniformly distributed and $t = ar$. In this case

$$\gamma_i(t) = \frac{q_i}{q_i + (1 - q_i)c}, \tag{5}$$

where $c = r_h/r_l > 1$ is a constant. This is clearly increasing in q_i. As a result, at any level of training at which there is uncertainty about resource access, members of a disadvantaged group will be thought to have a higher level of ability. If $q_1 < q_2$, it immediately follows that at a performance-maximizing allocation, any training level t contained in T_2 will also be contained in T_1.[6] This is true more generally.

Proposition 4. *If $q_1 < q_2$, then $T_2 \subseteq T_1$, with $T_2 = T_1$ if and only if $k \leq \hat{k}$.*

This means that a disadvantaged group will receive preferential treatment, not because of employers seeking to meet a representativeness target, but simply because the wish to maximize profits. Forcing the adoption of a pseudomeritocratic policy, therefore, will result in lower productivity. Of course the disadvantaged group could still be underrepresented on the whole.

4.3 A Special Case

The previous results regarding monotonicity and differential treatment can be illustrated by considering a spacial case in which $E(p_i|t)$ is increasing in t over the range $[0, t^*]$. A sufficient (but by no means necessary) condition for this is that a is uniformly distributed and training is linear in ability, since this implies that $\gamma_i(t)$ is constant. The following example, which assumes linear training and productivity functions, has this property, and the corresponding perfromance functions and training thresholds are shown in Figure 4.

Example 3. *Suppose the two groups are of identical size, $q_1 = 1/5$, $q_2 = 2/3$, $r_l = 1$, $r_h = 3/2$, $t = ar$, and $p = \beta a + (1 - \beta)t$, and $\beta = 4/5$. Then* $(\tilde{t}_1, \tilde{t}_2, t^*, \hat{t}_2, \hat{t}_1) = (0.76, 0.87, 1.00, 1.16, 1.31)$.

When $E(p_i|t)$ is increasing in t over the range $[0, t^*]$ for each i, we can obtain a sharp characterization of the performance-maximizing allocation in relation to elite capacity. For $k < \hat{k}$, we know that the allocation is pseudomeritocratic. For k slightly above $\hat{k}$, there exist $t \in (\tilde{t}_1, t^*)$ and $t'' \in (\hat{t}_2, \hat{t}_1)$ such that all individuals in $T_1 = [t, t^*] \cup [t'', \bar{t}]$ and $T_2 = [t'', \bar{t}]$ have expected productivity greater than all those outside this set, and the combined population of T_1 and T_2 is equal to k (see Figure 4). In this case only the the disadvantaged group faces non-monotonic selection. Similar reasoning implies that for even larger values of k, both groups face non-monotonic selection, and eventually, for k large enough, monotonic (but not pseudomeritocratic) selection results.

[6] The same argument applies if $t = a^{\delta} r^{1-\delta}$ for some $\delta \in (0, 1)$. In this case (5) holds with c replaced by $c^{(1-\delta)/\delta}$, so $\gamma_i(t)$ is again independent of t.

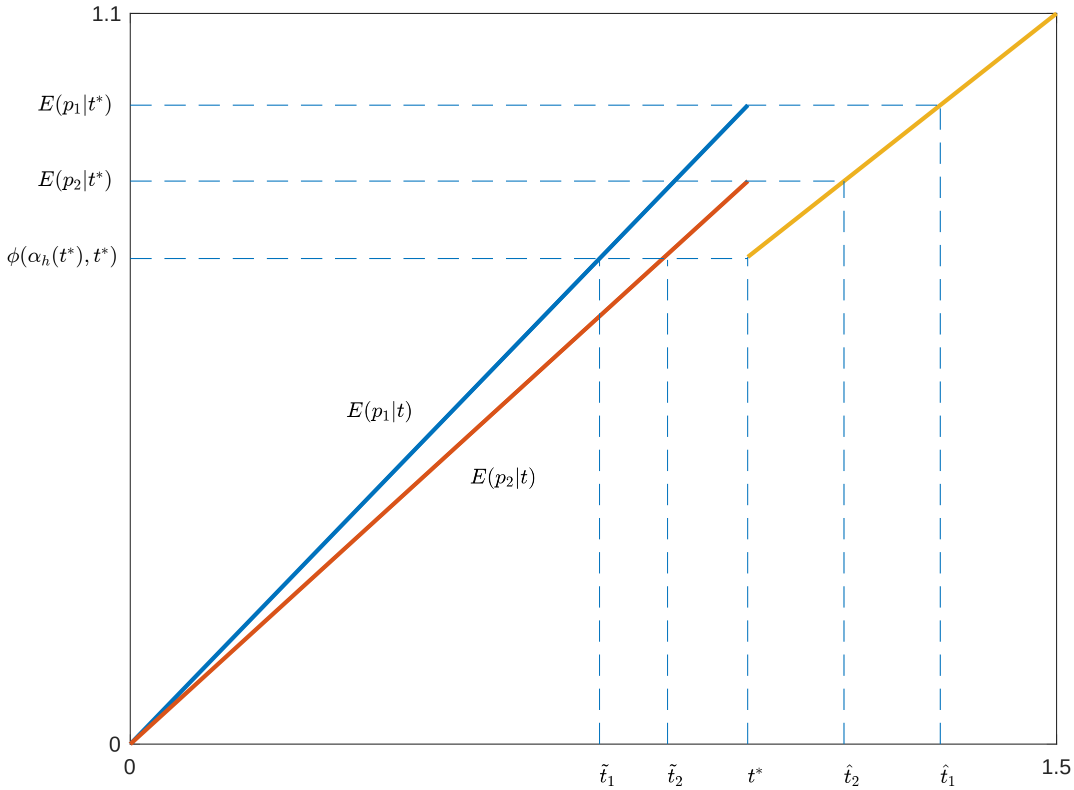

Figure 4: Uniform ability and linear training and performance functions.

Hence the performance-maximizing allocation must lie in one of four regimes, depending on elite capacity. If this sector is very small, then the allocation is pseudomeritocratic, only those with high resource access are selected, and a common standard is applied to the two groups. The disadvantaged group, having fewer individuals with high resource access, is underrepresented.

If elite capacity is somewhat larger, we have a regime in which those in the advantaged group are selected based on a single threshold, while those in the disadvantaged group face non-monotonic selection: some individuals below t^* are chosen, while some above this are not. This is because those with training slightly below t^* have higher ability in expectation if they belong to the disadvantaged group.

For even higher levels of elite capacity, both groups face nonmonotonic selection, but the standard applied to the disadvantaged group is less restrictive: some individuals with low levels of training are selected only if they belong to the disadvantaged group.

Finally, if elite capacity is large enough, the performance-maximizing allocation is again monotonic. But it is not pseudometritocratic: the threshold needed for selection from the disadvantaged group is lower than that needed for selection from the advantaged group.

4.4 Representation

Given any performance-maximizing allocation (T_1, T_2) at elite capacity k, let ρ_i be defined according to

$$\rho_i = \frac{1}{k}\left(q_i \int_{A_h(T_i)} dF(a) + (1-q_i)\int_{A_l(T_i)} dF(a)\right).$$

This is a measure of the degree to which group i is underrepresented in elite positions. It is easily verified that $\rho_1 = \rho_2 = 1$ if and only if the allocation is group egalitarian. The disadvantaged group is underrepresented (and the advantaged group overrepresented) if $\rho_1 < 1 < \rho_2$.

It is clear that the disadvantaged group will be underrepresented at any pseudomeritocratic allocation. To see this, consider any common threshold t applied to both groups, so $T_1 = T_2 = [t, \bar{t}]$. Then

$$k\rho_i = q_i(F(1) - F(\alpha_h(t))) + (1-q_i)(F(1) - F(\alpha_l(t))). \tag{6}$$

Since $\alpha_h(t) < \alpha_l(t)$, $q_1 < q_2$ implies $\rho_1 < \rho_2$. In particular, from Proposition **??**, the disadvantaged group is underrepresented at any performance-maximizing allocation when elite capacity is sufficiently small.

But this does not mean that the disadvantaged group will be underrepresented at *all* performance-maximizing allocations. The following example shows that in a performance-maximizing allocation the disadvantaged group can be *overrepresented*, despite having systematically lower resource access and hence lower levels of training on average.

Example 4. *Suppose the two groups are of identical size, $q_1 = 1/5$, $q_2 = 2/3$, $r_l = 1$, $r_h = 3/2$, $t = ar$, $p = \beta a + (1-\beta)t$, and $\beta = 5/6$. Then $t^* = 1$ and $\bar{t} = 3/2$. The performance-maximizing allocation at $k = 0.15$ is $T_1 = [t, t^*] \cup [t'', \bar{t}]$ and $T_2 = [t'', \bar{t}]$, where $t = 0.8815$ and $t'' = 1.1721$. At this allocation $\rho_1 = 1.03$.*

This example shows that the screening effect favoring a disadvantaged group can be strong enough to overcome the disadvantage itself.

Intuitively, the screening effect will be strong when productivity depends a great deal on ability relative to training. To explore this further, we consider a parametric specification for productivity given by

$$p = \phi(a, t) = \beta a + (1-\beta)t.$$

Here higher values of β correspond to a greater weight of ability relative to training. We know that the performance-maximizing allocation is pseudomeritocratic when $k < \hat{k}$, where $\hat{k}$ is defined in (4). Hence a higher value of $\hat{t}_1$ implies a lower value of $\hat{k}$, so nonmonotonic allocations and screening start to operate at lower levels of elite capacity. The following result shows how this threshold value of elite capacity varies with β.

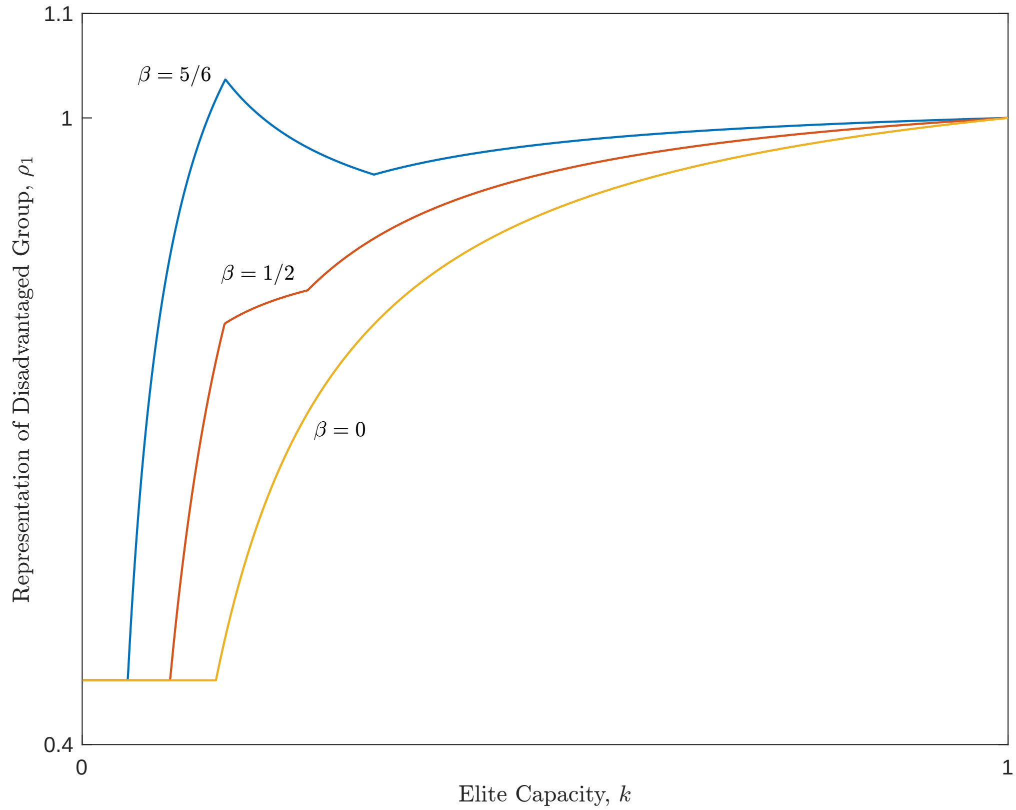

Figure 5: Underrepresentation of Disadvantaged Group for Various k and β.

Proposition 5. *The threshold level of elite capacity $\hat{k}$ below which the performance-maximizing allocation is pseudometritocratic is decreasing in β.*

To illustrate, consider the special case of a linear training function $t = \tau(a, r) = ar$ and uniformly distributed ability. In this case the degree of underrepresentation is constant in the regime with $k < \hat{k}$. Specifically, using (6) and the fact that $A_l(T_i)$ is empty in this regime, we obtain

$$\rho_i = \frac{q_i(1-\alpha_h(t))}{k} = \frac{q_i(1-\alpha_h(t))}{(s_1q_1+s_2q_2)(1-\alpha_h(t))} = \frac{q_i}{s_1q_1+s_2q_2},$$

which is independent of t. If $q_1 < q_2$ then $\rho_1 < 1 < \rho_2$, and the degree to which the disadvantaged group is underrepresented does not depend on k as long as this remains below $\hat{k}$.

Once k exceeds $\hat{k}$, screening starts to operate and the representation of the disadvantaged group grows. Proposition 5 states that this threshold is rising in β, and this can be seen in Figure 5 for a specific family of examples.[7]

The figure also illustrates the four regimes identified above. The lowest curve corresponds to the case where only training matters for productivity, in which case all performance-maximizing

[7]The figure is based on the same specifications as Example 4, but for additional values of β.

allocations are pseudomeritocratic and the disadvantaged group os always underrepresented. When ability also has an independent effect on productivity the disadvantaged group benefits from the screening effect, to a degree that is increasing in β. For high enough values of this parameter, there exists a range of values of k that result in overrepresentation of the disadvantaged.

The figure also shows that representation can vary non-monotonically with elite capacity. In particular, as the economy transitions from the second regime (in which only the disadvantaged group faces non-monotonic selection) to the third (in which both groups do), the representation of the disadvantaged group starts to fall when β is sufficiently high. Eventually, of course, equality of representation is ensured as k approaches 1, since all individuals secure elite positions.

4.5 Underinvestment

Finally, consider the incentive effects of non-monotonic policies. As in the discrete case considered in Section 3, when faced with such a policy, some individuals will have an incentive to attain levels lower levels of training than are within their reach, given their abilities. This could allow high resource individuals to signal that they have lower resource access, and hence higher ability, than would otherwise be the case. As a result, no individuals could be excluded from selection if anyone with lower training levels were selected.

Nevertheless, even in this case, the policy itself need not be monotonic. Indeed, for some parameter values, an equilibrium selection policy cannot be monotonic:

Proposition 6. *There exists an open set of parameter values such that the equilibrium selection policy is not monotonic.*

To see the reason for this, consider first the case of a single group. If elite capacity is such that a monotonic policy requires the threshold to be t^*, then expected performance can be increased by choosing a small interval of training levels above t^* and setting the selection probability for this set to zero. To respect the capacity constraint, a small interval below t^* can then be selected with positive probability. The newly rejected individuals will underinvest to pool with the newly accepted, but since the latter have strictly higher expected performance, overall performance will rise. This reasoning applies for all levels of elite capacity and all values of the population composition sufficiently close to these extreme cases.

If the equilibrium selection policy is not monotonic, then there must be some level of underinvestment: some set of individuals below t^* are admitted with positive probability while some above t^* are rejected, so the latter pool with the former. This pool now contains three types of individual: high and low resource types who don't underinvest, and high resource types who do.

As in the discrete case, if commitment to a selection policy is possible, then the same allocation of seats to candidates can be attained by choosing a monotonic policy that replicates the equilibrium. The reasoning is similar to that underlying the revelation principle: one simply assigns to each training level the probability of selection that the individuals in question would have attained in the equilibrium with underinvestment. This avoids the efficiency losses associated with underinvestment. But again, as in the discrete case, the pool quality is greater in the disadvantaged group (which has fewer high resource individuals to begin with), and the equilibrium policy will accordingly not be group-blind in general.

5 Conclusions

Meritocratic allocations are commonly understood to be those in which individuals best able to perform a given task are assigned to it. In a world in which these capacities cannot be directly observed and only imperfectly inferred through noisy signals such as test scores or years of completed education, we show that the nature of meritocratic allocations can be counter-intuitive and complex.

For instance, selection criteria need not be monotonic in signals of merit even within groups, disadvantaged groups may be favored in the selection process for reasons of productivity rather than representativeness, this screening effect can be strong enough to overwhelm the resource disadvantage itself in extreme cases, imposing pseudomeritocarcy can result in greater productivity losses than imposing group egalitarianism, and all these effects depend critically on the level of elite capacity. Through this exploration of the complex relationship between observed and actual merit, we see a central contribution of our paper as clarifying the concept of meritocracy and questioning policies based on its common interpretations.

Appendix

Proofs: Discrete Case

Proof of Proposition 1. Suppose $k < k_1$, the selection policy is pseudomeritocratic, and all individuals invest in their highest feasible training level. Since only those with training t_{hh} are selected, no other selection policy can result in higher expected performance. Since only (a_h, r_h) types can attain this training level, no individual has a profitable deviation.

Now suppose $k > k_2$, the selection policy is pseudomeritocratic, and all individuals invest in their highest feasible training level. Since only those with training t_{ll} are selected with probability below 1, no other selection policy can result in higher expected performance. Since only (a_l, r_l) types are excluded with positive probability, and they cannot invest at any level higher than t_{ll}, no individual has a profitable deviation.

If $k \in (k_1, k_2)$, a pseudomeritocratic policy must involve either

$$0 = \pi_1(t_{hl}) = \pi_2(t_{hl}) < \pi_1(t_{lh}) = \pi_2(t_{lh}),$$

or

$$0 < \pi_1(t_{hl}) = \pi_2(t_{hl}) < \pi_1(t_{lh}) = \pi_2(t_{lh}) = 1.$$

In either case, there is no incentive for underinvestment. Given the failure of performance monotonicity, performance can be increased by lowering the selection probability at t_{lh} and raising it at t_{hl}. Hence a pseudomeritocratic policy cannot arise in equilibrium. □

Proof of Proposition 2. Suppose $k \in (k_1, k_2)$ and $\lambda_1 > \mu$. We first consider the case $\lambda_2 > \mu$. Suppose that types (a_l, r_h) in both groups underinvest and choose t_{hl} while all other types invest at the highest feasible level. Then a selection policy with $\pi_i(t_{ll}) = \pi_i(t_{lh}) = 0$ and $\pi_i(t_{hh}) = 1$ is a best response if either $\pi_1(t_{hl}) = 1 > \pi_2(t_{hl})$, or $\pi_1(t_{hl}) > 0 = \pi_2(t_{hl})$. To see this, note that if $\pi_1(t_{hl}) < 1$ and $\pi_2(t_{hl}) > 0$ then expected performance can be raised by shifting probability from $\pi_2(t_{hl})$ to $\pi_1(t_{hl})$. Under the proposed selection policy, no individual can profitably deviate by changing her investment choice.

If $k \in (k_1, k_2)$ and $\lambda_1 > \mu > \lambda_2$, there is an equilibrium in which all types (a_l, r_h) in the disadvantaged group underinvest and choose t_{hl} while those of type (a_l, r_h) in the advantaged group distribute themselves across the training levels t_{hl} and t_{lh} in such a manner as to make the expected performance equal to μ at both training levels. Then a selection policy with $\pi_i(t_{ll}) = \pi_i(t_{lh}) = 0$ and $\pi_i(t_{hh}) = 1$ is a best response if either $\pi_1(t_{hl}) = 1 > \pi_2(t_{hl}) = \pi_2(t_{lh})$, or $\pi_1(t_{hl}) > 0 = \pi_2(t_{hl}) = \pi_2(t_{lh})$. To see this, note that if $\pi_1(t_{hl}) < 1$ and either $\pi_2(t_{hl})$ or $\pi_2(t_{lh}) > 0$ then expected performance can be raised by shifting probability from to $\pi_1(t_{hl})$ from either $\pi_2(t_{hl})$ or

$\pi_2(t_{lh})$. Under the proposed selection policy, no individual can profitably deviate by changing her investment choice.

Finally, suppose that $k \in (k_1, k_2)$ and $\mu > \lambda_1$. Then there is an equilibrium in which all types (a_l, r_h) in both groups distribute themselves across the training levels t_{hl} and t_{lh} in such a manner as to make the expected performance equal to μ at both training levels. Then a selection policy with $\pi_i(t_{ll}) = \pi_i(t_{lh}) = 0$ and $\pi_i(t_{hh}) = 1$ is a best response if $\pi_i(t_{hl}) = \pi_i(t_{lh})$. Under the proposed selection policy, no individual can profitably deviate by changing her investment choice. □

Proofs: Continuous Case

The following three preliminary results are useful for the proofs to follow.

Lemma 1. *For $t > t^*$, $E(p_1|t) = E(p_2|t)$, and $E(p_i|t)$ is increasing in t.*

Proof. For $t > t^*$, $\gamma_i(t) = 1$ for each i, and hence $E(p_i|t) = \phi(a_h(t), t)$. Since both ϕ and a_h are common to both groups and increasing in t, the claim follows. □

Lemma 2. *For each i, there exists a unique $\hat{t}_i \in (t^*, \bar{t})$ such that*

$$E(p_i|\hat{t}_i) = \max_{t \le t^*} E(p_i|t).$$

If $q_1 < q_2$ then $\hat{t}_2 < \hat{t}_1$.

Proof. Since $\phi(a, t)$ is increasing in both arguments and $t^* < \bar{t}$,

$$\max_{t \le t^*} E(p_i|t) < \phi(1, \bar{t}) = E(p_i|\bar{t}).$$

Since $\alpha_l(t^*) = 1 > \alpha_h(t^*)$, we have

$$E(p_i|t^*) = \gamma_i(t^*)\phi(\alpha_h(t^*), t^*) + (1 - \gamma_i(t^*))\phi(1, t^*) > \phi(\alpha_h(t^*), t^*).$$

For ε sufficiently small, $f(1) > 0$, imples $\gamma_i(t^*) < 1$ and hence

$$E(p_i|t^* + \varepsilon) = \phi(\alpha_h(t^* + \varepsilon), t^* + \varepsilon) < E(p_i|t^*) \le \max_{t \le t^*} E(p_i|t).$$

Recall that for $t > t^*$, $E(p_1|t) = E(p_2|t)$, and $E(p_i|t)$ is increasing in t. Hence there exists a unique $\hat{t}_i \in (t^*, \bar{t})$ such that $E(p_i|\hat{t}_i) = \max_{t \le t^*} E(p_i|t)$. The claim that $\hat{t}_2 < \hat{t}_1$ when $q_1 < q_2$ follows from Lemma 3. □

Lemma 3. *If $q_1 < q_2$, then $\tilde{t}_1 < \tilde{t}_2 < t^* < \hat{t}_2 < \hat{t}_1$.*

Proof. By definition

$$\max_{t \le t^*} E(p_i|t) = \phi(\alpha_h(\hat{t}_i), \hat{t}_i)$$

Since $E(p_1|t) > E(p_2|t)$ for all $t \le t^*$ and both $\alpha_h(t)$ and ϕ are increasing, we have $t^* < \hat{t}_2 < \hat{t}_1$. Similarly, by definition, $\hat{t}_i$ is the largest value of $t \le t^*$ such that

$$E(p_i|t) = \phi(\alpha_h(t^*), t^*).$$

Since $E(p_1|t) > E(p_2|t)$ for all $t \le t^*$, we have

$$E(p_1|\hat{t}_2) > E(p_2|\hat{t}_2) = \phi(\alpha_h(t^*), t^*)$$

which implies $E(p_1|t) > \phi(\alpha_h(t^*), t^*)$ for all $t \in [\hat{t}_2, t^*]$. Hence $\tilde{t}_1 < \tilde{t}_2 < t^*$. □

Proof of Proposition 3. If $k < \hat{k}$ then there exists $t' > \hat{t}$ such that

$$k = \sum_{i=1}^{2} s_i(1 - G_i(t')).$$

From Lemmas 1-2, for any $t'' > t' > \hat{t}$,

$$E(p_1|t'') = E(p_2|t'') > \max_{t<t'} E(p_i|t)$$

for each i. Hence the unique performance-maximizing allocation is given by $T_1 = T_2 = [t', 1]$, which is pseudomeritocratic.

Next consider the case $k \in (\hat{k}, \tilde{k})$. Suppose, by way of contradiction, that the performance-maximizing allocation is monotonic with thresholds t_1 and t_2. At least one of these thresholds must satisfy $t_i < \hat{t} = \max\{\hat{t}_1, \hat{t}_2\}$, otherwise we would have $k \le \hat{k}$. We separately consider two cases: $t_i < \hat{t}_i$, and $t_i < \hat{t}_j$ where $j \neq i$.

First consider the case $t_i < \hat{t}_i$. This implies $t_i < t^*$, since

$$E(p_i|t^*) > E(p_i|t)$$

for all $t \in (t^*, \hat{t}_i)$. But this in turn implies $t_i \le \tilde{t}_i$. If this were not the case, there would exist some $t \notin T_i$ with

$$E(p_i|t) > \phi(\alpha_h(t^*), t^*),$$

and hence

$$E(p_i|t) > E(p_i|t^* + \varepsilon)$$

for ε sufficiently small. Since $t^* + \varepsilon \in T_i$ and $t \notin T_i$, this is inconsistent with T_i being part of a performance-maximizing allocation.

Hence we have shown that $t_i \leq \tilde{t}_i$. But this implies that $t_j > \tilde{t}_j$, otherwise $k < \tilde{k}$ would be impossible. Now $t_j > \tilde{t}_j$ implies that there exists $t \notin T_j$ such that

$$E(p_j|t) > \phi(\alpha_h(t^*), t^*).$$

Hence expected productivity can be increased by replacing some individuals in T_i with training in $(t^*, t^* + \varepsilon)$ with an equal measure of those in group j with training in $(\tilde{t}_j, \tilde{t}_j + \varepsilon)$, for ε sufficiently small. This contradicts the hypothesis that (T_1, T_2) is a performance-maximizing allocation.

To complete the proof, consider the case $t_i < \hat{t}_j$ where $j \neq i$. In this case there exists $t \in T_i$ such that

$$E(p_i|t) < E(p_i|\hat{t}_j) = E(p_j|\hat{t}_j) = E(p_j|t^*).$$

This implies $t^* \in T_j$, and hence $t_j \leq \tilde{t}_j$ (following the same reasoning as above). As a result, $t_i > \tilde{t}_i$, otherwise $k < \tilde{k}$ would be impossible. Hence there exists $t \notin T_i$ such that

$$E(p_i|t) > \phi(\alpha_h(t^*), t^*).$$

Hence expected productivity can be increased by replacing some individuals in T_j with training in $(t^*, t^* + \varepsilon)$ with an equal measure of those in group i with training in $(\tilde{t}_i, \tilde{t}_i + \varepsilon)$, for ε sufficiently small. This contradicts the hypothesis that (T_1, T_2) is a performance-maximizing allocation.

To prove the last claim, note that since ϕ is increasing in both arguments, the following holds for each $t > 0$ and each i:

$$E(p_i|0) < E(p_i|t).$$

That is, expected productivity is minimized at the lowest attainable level of training, since this also involves the lowest level of ability. Hence there exists $t' > 0$ such that, for all $t'' \leq t'$ and $t > t'$, and each group i,

$$E(p_i|t'') < E(p_i|t).$$

In this case, regardless of group membership, if an individual with training t' secures an elite position, so must all individuals with $t > t'$. And if an individual with training t' fails to secure an elite position, those with $t < t'$ must also fail to do so. For any given t', if k is sufficiently large, then at least some individuals with $t < t'$ in each group must be assigned to elite positions. Hence the allocation is monotonic. □

Proof of Proposition 4. Let $H(q, t)$ be defined as

$$H(q,t) = \frac{qf(\alpha_h(t))a_h'(t)}{qf(\alpha_h(t))a_h'(t) + (1-q)f(\alpha_l(t))\alpha_l'(t)},$$

and note that

$$\frac{\partial H(q,t)}{\partial q} = \frac{f(\alpha_h(t))f(\alpha_l(t))a_h'(t)a_l'(t)}{(qf(\alpha_h(t))a_h'(t) + (1-q)f(\alpha_l(t))\alpha_l'(t))^2} > 0.$$

Since $\gamma_i(t) = H(q_i, t)$, $q_1 < q_2$ implies $\gamma_1(t) < \gamma_2(t)$ for $t \in [\tau(0, r_h), t^*]$. Recall that expected productivity conditional on training is

$$E(p_i|t) = \gamma_i(t)\phi(\alpha_h(t), t) + (1 - \gamma_i(t))\phi(\alpha_l(t), t).$$

Since $\alpha_h(t) < \alpha_l(t)$, we obtain

$$q_1 < q_2 \implies E(p_1|t) > E(p_2|t)$$

at any $t \in [\tau(0, r_h), t^*]$.

If $k > \hat{k}$, then there must be at least some $t \in [0, t^*)$ that is contained in T_1. If there is no such t in T_2 then clearly $T_1 \neq T_2$ as claimed. Suppose, instead, that there exists some $t \in [0, t^*)$ that is contained in T_2. Since $E(p_1|t) > E(p_2|t)$, this must also be contained in t_1. Furthermore, since $k < 1$, there must be some such $t \in T_2$ such that no neighborhood of t is contained in T_2. However since $E(p_1|t) > E(p_2|t)$, the set $(t - \varepsilon, t + \varepsilon)$ must be contained in T_1 for ε sufficiently small. Hence $T_1 \neq T_2$ as claimed. □

Proof of Proposition 5. Fix any group i, set $\gamma = \gamma_i(t^*)$, and define $H(\beta, t)$ as follows

$$\begin{aligned}
H(\beta, t) &= \gamma\phi(\alpha_h(t^*), t^*) + (1 - \gamma)\phi(\alpha_l(t^*), t^*) - \phi(\alpha_h(t), t) \\
&= \gamma(\beta\alpha_h(t^*) + (1 - \beta)t^*) + (1 - \gamma)(\beta\alpha_l(t^*) + (1 - \beta)t^*) - (\beta\alpha_h(t) + (1 - \beta)t) \\
&= \beta(\gamma\alpha_h(t^*) + (1 - \gamma)\alpha_l(t^*)) + (1 - \beta)t^* - (\beta\alpha_h(t) + (1 - \beta)t).
\end{aligned}$$

Note that $\hat{t}_i$ is a solution for t to $H(\beta, t) = 0$ and that $\partial H/\partial t$ is everywhere nonzero. This implicitly defines a function $\hat{t}_i(\beta)$. Using the Implicit Function Theorem, we obtain

$$\frac{d\hat{t}_i}{d\beta} = -\frac{\partial H/\partial \beta}{\partial H/\partial t} = \frac{\gamma\alpha_h(t^*) + (1 - \gamma)\alpha_l(t^*) - \alpha_h(\hat{t}_i) + \hat{t}_i - t^*}{\beta a_h'(\hat{t}_i) + (1 - \beta)}$$

Since $\hat{t}_i > t^*$, the above expression is also positive provided that

$$\gamma\alpha_h(t^*) + (1 - \gamma)\alpha_l(t^*) > \alpha_h(\hat{t}_i). \tag{7}$$

To see that this must be the case, note that by definition,

$$\gamma(\beta\alpha_h(t^*) + (1 - \beta)t^*) + (1 - \gamma)(\beta\alpha_l(t^*) + (1 - \beta)t^*) = \beta\alpha_h(\hat{t}_i) + (1 - \beta)\hat{t}_i.$$

Hence

$$\beta(\gamma\alpha_h(t^*) + (1 - \gamma)\alpha_l(t^*) - \alpha_h(\hat{t}_i)) = (1 - \beta)(\hat{t}_i - t^*) > 0,$$

which implies (7). Hence $\hat{t}_i$ is increasing in β. The claim then follows from the definition of $\hat{k}$. □

Proof of Proposition 6. Define $k^* \in (0,1)$ as follows:

$$k^* = \sum_{i=1}^{2} s_i(1 - G_i(t^*)).$$

Suppose first that $k = k^*$ and $s_2 = 0$ (so there is just one group). If the equilibrium selection policy is monotonic, then $t_1 = t^*$. Consider the following perturbation of the selection policy: select those group 1 individuals in $[t^* + \varepsilon, \bar{t}]$ with probability 1, and those in $[t^* - \varepsilon, t^*]$ with probability $r < 1$. The pair (ε, r) is chosen to ensure that the capacity constraint continues to be satisfied, conditional on all group 1 individuals in $(t^*, t^* + \varepsilon)$ choosing training levels in $[t^* - \varepsilon, t^*]$, and no other underinvestment. This will be possible if ε is sufficiently small. Also, for ε sufficiently small, the perturbed policy increases expected performance, since it replaces some individuals slightly above t^* with some who are slightly below t^*. The performance of those newly accepted is strictly greater than those newly rejected, so the perturbed policy is performance increasing. Clearly a monotonic selection policy will fail to be optimal for k sufficiently close to k^*.

Now suppose that $s_2 > 0$ and $k = k^*$. Under any monotonic policy, we must have $t_1 = t^* + \eta$, where $\eta \geq 0$. By choosing s_2 sufficiently small, η can be made arbitrarily close to zero. The above reasoning then applies: performance can be increased by cutting the probability of selection for a small interval above t_1 to zero, and admitting a small interval below t^* with probability chosen to respect the capacity constraint (holding constant the selection policy for group 2).

This shows that a non-monotonic policy exists that is strictly superior to the pseudomeritocratic policy at elite capacity k^*. The claim then follows from the continuity of expected performance in elite capacity. □

Robustness

The previous arguments have been made under the assumption that there are some levels of training that are unattainable for those with low resource access. This seems reasonable, and considerably simplifies the argument, but is not necessary for the identified effects to arise.

To see this, consider the following example. Let $g_h(t) = f(\alpha_h(t))a'_h(t)$ and $g_l(t) = f(\alpha_l(t))a'_l(t)$ denote the density functions for training conditional on high and low resource access, and suppose that these are given by beta distributions with shape parameters $(2, 5)$ and $(5, 2)$ respectively. The underlying distribution of ability $f(a)$, the two levels of resource access r_l and r_h, and the mapping $\tau(a, t)$ from ability and resources to training are all implicit here, and not uniquely pinned down. Suppose that $p = a^{\omega} t^{1-\omega}$ and that $\omega = 0.8$.

The mean levels of training in this case are 2/7 for those with low resource access and 5/7 for those with high resource access, and the corresponding modes are at 0.2 and 0.8. As long as

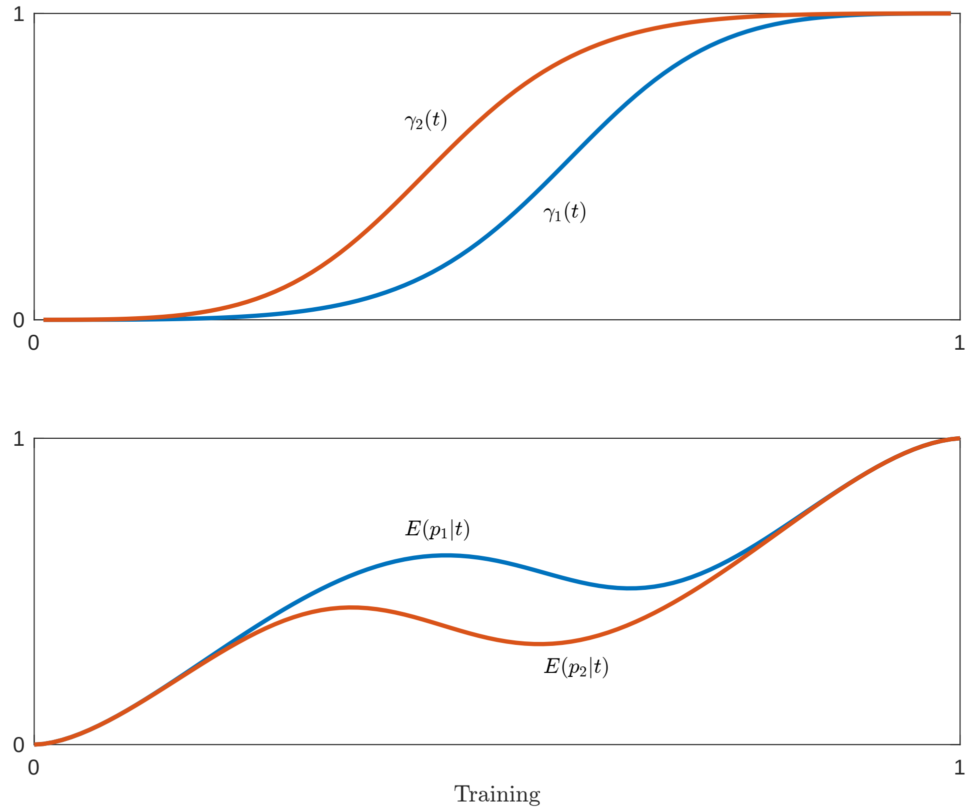

Figure 6: Likelihood of High Resource Access and Expected Productivity by Group.

q_i is not too extreme, at low levels of training an individual is very likely to have low resource access. Likewise, at high training levels, one is very likely to be dealing with an individual having high resource access. Nevertheless, all levels of training are reachable at both resource levels so uncertainty is never resolved.

Now suppose that the groups have composition $q_1 = 0.3$ and $q_2 = 0.7$ respectively. In this case expected productivity conditional on training will differ by group. At any given training level, an individual's expected level of ability will be higher if she belongs to the disadvantaged group, and this effect is especially strong at intermediate levels of training.

Figure 6 shows how $\gamma_i(t)$ and $E(p_i|t)$ vary with t in this example. The likelihood that one is dealing with someone having high resource access is increasing in t for both groups, but at different rates. At intermediate values of training, someone in a disadvantaged group is considerably more likely to have low resource access, and hence much more likely to have higher productivity. This productivity gap may be seen in the bottom panel of the figure.

The figure also shows that expected productivity does not vary monotonically with training in either group. The performance-maximizing allocation will be monotonic if elite capacity is sufficiently high or low, but at intermediate levels of capacity one or both groups will face non-monotonic selection criteria. Furthermore, equalization of accepted productivity among the marginal accepted types implies that T_2 will be a proper subset of T_1 at all values of k, which is even stronger than the claim in Proposition 4. Hence none of our main qualitative claims depend on the existence of discontinuities in expected productivity conditional on training, which arise when some training levels are out of reach for those with low resource access.

References

Dennis J. Aigner and Glen G. Cain. Statistical theories of discrimination in labor markets. Industrial and Labor Relations Review, pages 175–187, 1977.

Kenneth J. Arrow. The theory of discrimination. In Orley Ashenfelter and Albert Rees, editors, Discrimination in Labor Markets. Princeton, NJ: Princeton University Press, 1973.

Dario Cestau, Dennis Epple, and Holger Sieg. Admitting students to selective education programs: Merit, profiling, and affirmative action. Journal of Political Economy, 125(3):761–797, 2017.

Jimmy Chan and Erik Eyster. Does banning affirmative action lower college student quality? American Economic Review, 93(3):858–872, 2003.

Stephen Coate and Glenn C. Loury. Will affirmative-action policies eliminate negative stereotypes? The American Economic Review, 83(5):1220–1240, 1993.

Bradford Cornell and Ivo Welch. Culture, information, and screening discrimination. Journal of Political Economy, pages 542–571, 1996.

Steven N. Durlauf. Affirmative action, meritocracy, and efficiency. Politics, Philosophy & Economics, 7(2):131–158, 2008.

Roland G. Fryer and Glenn C. Loury. Valuing diversity. Journal of Political Economy, 121(4): 747–774, 2013.

Roland G. Fryer, Glenn C. Loury, and Tolga Yuret. An economic analysis of color-blind affirmative action. Journal of Law, Economics, and Organization, 24(2):319–355, 2008.

Glenn C. Loury. Intergenerational transfers and the distribution of earnings. Econometrica, 49(4): 843–867, 1981.

Edmund S. Phelps. The statistical theory of racism and sexism. American Economic Review, 62 (4):659–661, 1972.

Debraj Ray and Rajiv Sethi. A remark on colorblind affirmative action. Journal of Public Economic Theory, 12(3):399–406, 2010.

John E Roemer. Equality of opportunity. Harvard University Press, 2009.

TM Scanlon. Why Does Inequality Matter? Oxford University Press, 2018.