Bias and Diveristy III: Valuing Diversity

Professor Glenn Loury
Brown University
Department of Economics

10/30/17

These notes are based off a presentation by Professor Glenn Loury (Brown, Department of Economics), for the section on "Bias and Diversity" in the Mechanism Design for Social Good Reading Group. The notes are taken by members of the reading group with some figures and texts taken from the accompanying paper, Valuing Diversity [2]. Questions and comments from reading group members during the presentation are labeled as such. Please contact the reading group organizers with any questions or comments.

- This paper looks at the economics of diversity-enhancing policies. We develop a simple two-stage model: in the first stage, heterogeneous agents with different costs of skill acquisition decide whether or not to invest in skills. In the second stage, these agents (with different productivities) enter a competitive market where they can hold slots, which are productive opportunities, that enable them to use their skills. One group is disadvantaged in the sense that acquisition of skill has higher cost. We want to enhance diversity by increasing representation of the disadvantaged group in this setting.

- Main question: how do we design interventions in this process of competitive allocation of opportunity so as to promote the participation by disadvantaged population group in these activities that are being rationed?

- The supposition is that regulation/policy makers have decided to intervene with the goal of increasing and enhancing the presence of some population group in an activity where opportunities for pursuing it are scarce.

- Top-level division: where in the economic life-cycle should such an intervention take place to accomplish the above goal? This is a *timing* concern.
  - Before productivities of agents are determined: enhance *developmental opportunities*
  - After productivities are determined: ration opportunities to favor some population

- *Question:* Is there an assumption that after a fixed point in time people stop developing skills?

- *Answer:* Yes, it's a two-stage process: in the first stage agents acquire skills, and in the second stage opportunities are rationed. This may not be so clear in real life. As an example, think of faculty appointments: the first stage is for students, where they develop skills. The second stage is at the time of hiring, when faculty opportunities are rationed. In practice, the gradations will not be a binary process of development then opportunity, but for the purposes of the model, we consider this discrete binary timing. One could think about what would happen if you weakened this assumption; however, conceptually, we do want to distinguish between the development margin and the opportunity margin, as people do seem to think differently about these from an ethical perspective.

- *Comment:* It seems like some sort of function that is not as discretized might better capture the real world. For instance, something that captures that opportunities give you more of a chance to develop further skills. We can weaken this before-or-after analysis and see how the results carry over. But, at the very least, we do have to differentiate between the above two margins conceptually.

- We could ask, from an efficiency point of view, where you want to put your thumb on the scale. "Laissez-faire," which is the policy of not intervening, is the most efficient way to allocate the resources in our world. People are making investment decisions that are costly to them but enhance their productivity. And then, someone is occupying positions where they can produce, and others are not occupying that position. These scarce opportunities are being rationed. This describes two margins where economically relevant activity is taking place. Because one population is disadvantaged, that group will invest in skills at a lower rate and will be underrepresented out of those who receive opportunities.

  On the other hand, opportunities (or lack there of) allow people to practice (or not) the skills they acquire. This opportunity is also scarce, so you want the right people (the most skilled people) to make use of that opportunity. If you do not intervene in any way, then you would get the most efficient solution where the scarce opportunities go to the most skilled people and the investments to acquire skills will be made by those who find it the least costly to make those investments, i.e., the advantaged group. Therefore, the disadvantaged group will be underrepresented. One wants to correct for this natural outcome, in line with our goal above. Any intervention will require some loss of efficiency, where efficiency is defined as the solution which maximizes net social surplus. The question is then *what is the least costly way to increase representation of a disadvantaged group?*

- When we depart from laissez-faire, do we want to do this at the point of investment or the point of allocation? This question of the development margin vs. the opportunity margin is of great importance.

- *Question:* When we talk about efficiency, do we mean instantaneous efficiency or efficiency in some sort of long-term dynamical way?

- *Answer:* Efficiency is social welfare, and the model has no dynamic component (the distribution of investment costs is exogenous): there is no evolution of investment distributions. We're concerned here with instantaneous efficiency (after both stages) in this paper.

- *Question:* Is this different from misallocation of talent, since the cost distributions are exogenous?

- *Answer:* Yes, there is no misallocation of talent. With laissez-faire, opportunities go to the "most talented people." The only reason this may be violated is if we intentionally depart from this to improve representation.

- Another dimension to consider in addition to timing is *visibility*. To intervene, I can do so either directly or indirectly.

  - Directly intervening may result in different *cutoffs* for each group, where people above the cutoff get an opportunity and people below don't. This means that two people with the same score but from different groups may have different outcomes. This may be objected to because it explicitly uses racial information to make decisions.
  - Indirectly means group-blind. "Color-blind affirmative action" treats people the same way regardless of race, but it's calibrated to boost one group. For example, people are uniformly distributed on the unit interval in terms of their productivity. Suppose we can only accommodate for 20 percent of the people in the slots. Suppose we have group A being more favorably distributed in the unit interval than group B. If we apply the common cutoff of 0.8, then we will have B being underrepresented. We instead apply the following rule:

    * Everyone with productivity above 0.9 gets a slot.
    * For everyone else, we hold a lottery and give everyone a uniform 1/9 chance of getting a slot.

    This does not explicitly use race, but is calibrated so as to produce higher representation of the disadvantaged group. It will be costly with respect to efficiency, (i.e. someone with a productivity of 0.2 getting a slot with positive probability). But, we are able to increase diversity and respects the information constraint and enhances representation.

- *Question:* What is the objective of the intervention?

- *Answer:* We have groups A and B, characterized by an investment cost distribution. If the skill is acquired, an individual's productivity is drawn from a distribution that is stochastically bigger than if the skill is not acquired. Agents decide whether or not to invest, and then they either do or do not occupy a slot which gives an opportunity to produce. If they don't get a slot, they get nothing, and if they do, they get their realized productivity. This will produce net output (productivity minus investment costs) and an allocation of people: some number of group A members get slots and some number of B members get slots. The goal of the intervention is to narrow the gap between the A's and B's – moving toward parity. If we don't intervene, the proportion of B's given opportunities would be some ρ_B^*, which is less than θ, the proportion of the general population given opportunities. We're given a number $\rho_B \in [\rho_B^*, \theta]$, which is the target rate at which B's are given opportunities. The goal is to maximize social welfare subject to the constraint that we meet the ρ_B rate for B's. (We assume that $\rho_B > \rho_B^*$, and we assume that this ρ_B is given as a policy target.)

- *Question:* What's the space of color-blind allocation rules we consider? Can we randomize between arbitrary productivity levels?
- *Answer:* Yes, in general, a way of allocating slots is to have some function $a(v)$ over productivities v where $a(v)$ is the probability that someone with productivity v gets a slot. The most efficient allocation will be a step function $a(v)$: 0 below some threshold and 1 otherwise. We consider non-decreasing functions to make it incentive-compatible. This is an infinite-dimensional linear program, which we can solve.
- Result: Under a certain assumption, the optimal color-blind allocation has the property that for some threshold, everyone above the threshold gets a slot and everyone below gets a slot with constant probability. The parameters are determined by the exogenous capacity and representation constraints.
- *Question:* Are there examples of policies in the world where people come close to explicitly doing this randomization (as opposed to implicitly randomizing but using other attributes)?
- *Answer:* All the ones I can think of are things like percentage plans (where the top x% of a high school gets university admission) are of the form where non-identity factors that are correlated with identity are used. However, something like recommendation letters, which are very noisy, could be a form of randomization. The randomization is a device in this model, and is perhaps an abstract way of representing what actually happens in practice.
- *Question:* Is there some intuition as to why the optimal color-blind policy has the property described above?
- *Answer:* Figure 1 gives some intuition for a swapping argument, where any policy that doesn't follow this structure can be transformed to one that does by only making it better. The x-axis is productivity, and the y-axis is the allocation rule. In this figure μ is productivity and α is probability of getting a slot. This figure shows do slots get allocated as a function of reported productivity. The two policies depicted are the one asserted to be optimal, which is the step function, and any alternative policy. The pluses and minuses show that any alternative policy must shift mass towards the middle of the distribution, and we can show that any shift away from the tails makes things worse. The assumption needed to make this true is shown in Figure 2. The plot shows the ratio

$$\frac{\text{rate at which people in group } B \text{ exhibit productivity } \mu}{\text{rate at which people in group } A \text{ exhibit productivity } \mu}. \tag{1}$$

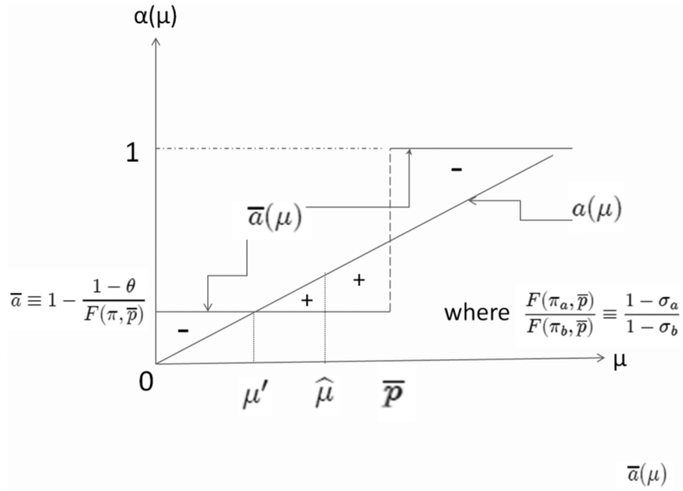

FIG. 3.—Uniqueness of optimal stage 2 policy: any alternative feasible policy $a(\mu)$ must shift weight to the middle from the extremes of the range of μ, relative to $\bar{a}(\mu)$.

Figure 1: Swapping argument

This function has to be decreasing by the model assumptions (group B is disadvantaged relative to A). The additional assumption we make is that it's convex. The intuition

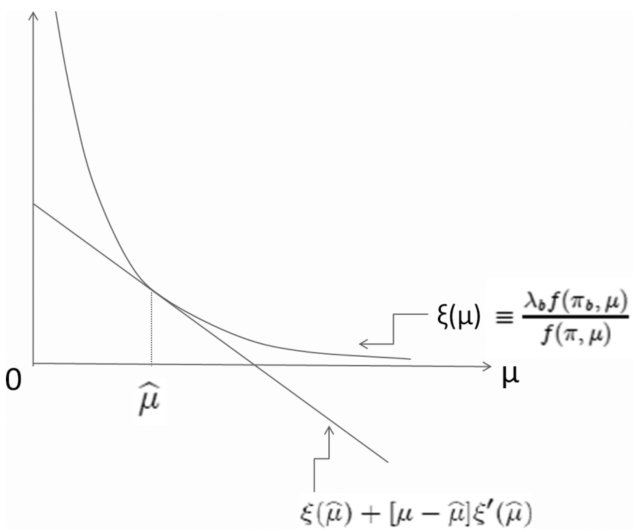

Figure 2: Convexity assumption

for why this is the case is that any shift away from the optimal solution must shift more away from the upper tail since fewer group B people are at the upper tail.

- Another way of seeing this is that this is a linear program, so the solution has to be at a vertex of the polytope of feasible set, and the given solution cannot be written as a convex combination of any other solutions. This doesn't show it's optimal, but confirms that it's at least a plausible solution.

- *Question:* Does the assumption imply the following: as I increase the threshold for a guaranteed slot, this monotonically increases representation until, when this threshold goes to 1, parity is achieved? And is this equivalent to the assumption?

- *Answer:* Yes to the first part, but the assumption is stronger than this condition. In fact, this condition is a consequence of the assumption that we have a monotone likelihood ratio property in the conditional distribution of productivity given investment and the distribution of investment cost given group. The convexity assumption is an additional constraint beyond this.

- We are interested in group blindness and policy. We have more results:

  - If you don't require blindness, then the optimal intervention occurs at assignment of opportunities and not at the development margin where people are deciding investment.

  - If you require blindness, you might also want to intervene at the development margin (in addition to the intervention described above). This means either subsidizing or taxing the acquisition of skills. Whether you subsidize or tax depends on whether group B is better represented at the development or assignment margin.

    * Development margin: people who are indifferent between acquiring or not acquiring skills (level of cost).
    * Assignment margin: people who are just indifferent between getting a slot for sure or not (level of productivity).
    * If B's are more represented at development than assignment margin, optimal policy will subsidize acquisition of skill. Otherwise, optimal policy will *tax* skill acquisition.

- *Question:* If the optimal policy has a step-function property, does this tell me something about the convexity of the conditional probability ratio?

- *Answer:* Suppose we didn't impose convexity. Then, qualitatively, optimal solution has to be step function with no more steps than constraints (2 in our case). See [1].

- *Question:* What is one assumption you'd like to relax in this model?

- *Answer:* Earlier we talked about whether there were any dynamics to this model. This model doesn't address where the disadvantage comes from and how it can change. With some policies (i.e. reservation policies in India), the assumption seems to be that the need for them will change over time.

  Furthermore, how we determine "productivity" is a point of discussion – there isn't just one observable number that determines productivity. It's not clear that productivity is an unambiguous thing. What if we need some mix of characteristics? Suppose you have two dimensions, and maybe group B is better with respect to one characteristic but not another. This might affect how we think about this problem.

References

[1] E.J. Anderson and P. Nash. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. Wiley-Interscience Series in Discrete Mathematics and Optimization Series. Books on Demand, 1987.

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