Allocating Opportunities in a Dynamic Model of Intergenerational Mobility (Extended Abstract)

Hoda Heidari[1*], **Jon Kleinberg**[2]
[1]Carnegie Mellon University
[2]Cornell University
hheidari@cs.cmu.edu, kleinberg@cs.cornell.edu

Abstract

Opportunities such as higher education can promote *intergenerational mobility*, leading individuals to achieve levels of socioeconomic status above that of their parents. In this work, which is an extended abstract of a longer paper in the proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, we develop a dynamic model for allocating such opportunities in a society that exhibits bottlenecks in mobility; the problem of optimal allocation reflects a trade-off between the benefits conferred by the opportunities in the current generation and the potential to elevate the socioeconomic status of recipients, shaping the composition of future generations in ways that can benefit further from the opportunities. We show how optimal allocations in our model arise as solutions to continuous optimization problems over multiple generations, and we find in general that these optimal solutions can favor recipients of low socioeconomic status over slightly higher-performing individuals of high socioeconomic status — a form of socioeconomic affirmative action that the society in our model discovers in the pursuit of purely payoff-maximizing goals. We characterize how the structure of the model can lead to either temporary or persistent affirmative action, and we consider extensions of the model with more complex processes modulating the movement between different levels of socioeconomic status.

1 Introduction

Intergenerational mobility — the extent to which an individual's socioeconomic status differs from the status of their prior generations of family members — has emerged as a central notion in our understanding of inequality. A large amount of empirical work has gone into estimating the extent of mobility for different subsets of society; while many of the effects are complex and challenging to measure, two broad and fairly robust principles emerge from this work. First, socioeconomic status is persistent across generations: an individual's socioeconomic status is strongly dependent on parental status. As Lee and Solon [2009] write in the opening to their survey of this topic, *"Over the past two decades, a large body of research has documented that the intergenerational transmission of economic status in the United States is much stronger than earlier sociological and economic analyses had suggested"*. Second, certain types of opportunities can serve as strong catalysts for socioeconomic mobility; a canonical example is higher education, which has the potential to raise an individual's socioeconomic status (and, by the previous principle, that of their current or future children as well). As Chetty et al. [2014] write, *"The fact that the college attendance is a good proxy for income mobility is intuitive given the strong association between higher education and subsequent earnings"*.

An important question from a social planning perspective is thus the choice of policy for allocating opportunities to people of different levels of socioeconomic status. (Again, we can think of access to higher education as a running example in this discussion.) Many goals can motivate the choice of policy, including the reduction of socioeconomic inequality and the prioritization of opportunities to those most in need. Such goals are often viewed as operating in tension with the aim of maximizing the achievable payoff from the available opportunities, which would seem to suggest targeting the opportunities based only on the anticipated performance of the recipient, not their socioeconomic status. In this view, society is implicitly being asked to choose between these goals; this consideration forms a central ingredient in the informal discourse and debate around the allocation of opportunity. But through all of this, a challenging question remains: to what extent is the tension between these goals genuine, and to what extent can they be viewed as at least partially in alignment?

A large body of work in economics compares various allocation policies in terms of the above seemingly-competing criteria — typically in simplified settings in which only two generations are considered. The literature includes seminal work by Nobel Laureate Garry Becker with Nigel Tomes [1986] and by Glenn Loury [1981]. In multigenerational settings, however, deriving the optimal policy becomes exceedingly challenging, and it has been highlighted as a class of open questions in this literature. For example, in his work on models of college admissions and intergenerational mobility, Durlauf [2008] notes: *" A college admissions rule*

*Contact Author

has intergenerational effects because it not only influences the human capital of the next generation of adults, but also affects the initial human capital of the generation after next. [...] Efficiency in student allocation [in this case] is far more complicated than before. I am unaware of any simple way of describing efficiency conditions for college assignment rules analogous to [the above setting]." In this work, we address this challenge and the associated open questions concerning the behavior of multigenerational models. A key ingredient in our progress on these questions is the development of methods for working with a class of Markov Decision Processes that operate over continuous states and continuous actions. Our analysis of multigenerational models enables us to investigate the apparent tension between efficiency and fairness considerations in allocating opportunities.

We work with a simple mathematical model representing a purely payoff-maximizing society, operating over multiple generations. Our model is grounded in the types of models proposed in economic theory work on these problems [Durlauf, 2008]. The society must decide how to allocate opportunities in each generation across a population heterogeneous in its socioeconomic status. The payoff to the society is the total performance of everyone who receives the opportunities, summed (with discounting) over all generations. Although the set-up of the model is highly streamlined, the analysis of the model becomes quite subtle since society must solve a continuous-valued dynamic programming problem over multiple generations.

What we find from the model is that the optimal solution will in general tend to offer opportunities to individuals of lower socioeconomic status over comparable individuals of higher socioeconomic status, even when these competing individuals are predicted to have a slightly better performance from receiving the opportunity. This is not arising because the optimal solution has any a priori interest in reducing socioeconomic inequality (although such goals are important in their own right [Forde-Mazrui, 2004]); rather it is strictly trying to maximize payoff over multiple generations. But given two individuals of equal predicted performance, the one with lower socioeconomic status confers an added benefit to the payoff function: their success would grow the size of the socioeconomically advantaged class, resulting in higher payoffs in future generations. Because the difference in payoff contributions between these two individuals is strictly positive, the same decision would be optimal even if the individual of lower socioeconomic status had a slightly lower predicted performance from receiving the opportunity. The optimal solution should still favor the candidate with lower status in this case.

In other words, the society in this model discovers a form of socioeconomic affirmative action in allocating opportunities, based purely on payoff-maximizing motives. The model thus offers a view of a system in which reducing inequality is compatible with direct payoff maximization. In this sense, our results belong to a genre of analyses (popularized by Page [2008] and others) asserting that policies and interventions that we think of as motivated by equity concerns, can also be motivated by purely performance-maximizing considerations: *even if* society only cares about performance, not equity, it should still (at least in the underlying models) undertake these policies. In addition to providing a purely utilitarian motivation for socioeconomic affirmative action, our model provides novel insights regarding the shape and extent of effective affirmative action policies by specifying the way in which criteria for receiving the opportunity should be adjusted based on socioeconomic status to maximize society's performance across multiple generations.

2 Overview of the Model

We consider a population, represented by a continuum of agents, each belonging to one of the two socioeconomic status groups: group D (disadvantaged), consisting of a ϕ_0 fraction of the population, and group A (advantaged), consisting of a $\phi_1 = 1 - \phi_0$ fraction of the population. Each agent i (from either group) has an ability a_i drawn uniformly at random from the interval $[0, 1]$,

Society has the ability to offer an opportunity to an α fraction of the population. Note that the parameter α specifies the inherent limitation on the amount of opportunities available. Since opportunities are limited, the society has to wrestle with the question of how to allocate them. An individual i in group D who is offered the opportunity has a probability σa_i of succeeding at it, for a parameter $0 < \sigma < 1$. An individual i in group A who is offered the opportunity has a probability $\sigma a_i + \tau$ of succeeding at it, for the same σ and an additional parameter $0 < \tau \leq 1 - \sigma$ reflecting the advantage. We will refer to the above quantities as the *success probabilities* of the agents. Success probabilities reflect various levels of performance when agents are offered the opportunity.

Anyone in group D who is offered the opportunity and succeeds at it moves up to group A. Each individual is then replaced by one offspring of the same socioeconomic status in the next generation and the process continues. In the general form of the model, there is also some probability that an individual's offspring does not perfectly inherit their socioeconomic status. The payoff to society is the number of individuals who succeed at the opportunity summed over all generations, with the generation t steps into future multiplied by γ^t for a discount factor $0 < \gamma < 1$.

In any given generation, society's policy will consist of a threshold, θ_D, for group D and a (possibly different) threshold, θ_A, for group A: the opportunity is given to every individual whose success probability is above the threshold for their group. The optimal policy is given by a dynamic program over the continuous set of all possible choices for the population composition (ϕ_0, ϕ_1) as state variables.

Our model defines a class of Markov Decision Processes (MDPs) that operate over continuous states and continuous actions. In our model, states correspond to different composition of the population (specified by ϕ_D), and actions at each state correspond to the threshold θ_D applied to group D in that state (note that both the percentage and threshold for A is determined automatically given phi_D and θ_D.) We are interested in understanding whether and to what extend the optimal policy in these MDPs use affirmative action.

3 Summary of Results

We solve the most basic version of the above MDPs analytically. We also computationally solve more complex versions of the model by discretizing the state space, then applying standard dynamic programming solutions for finite decision processes.

If the problem of allocating the opportunity only spanned a single generation, then the payoff-maximizing policy would use the same threshold for both groups. But given the discounting sum over multiple generations, we find that society's optimal policy can, in general, use a lower threshold for group D than for group A. The difference in thresholds is a form of socioeconomic affirmative action, and it arises due to the intuition discussed above: boosting the number of individuals from group D who receive the opportunity will increase the number of available candidates from group A in future generations, each of whom provides a (discounted) payoff in future generations via their enhanced performance. Finding the correct trade-off in allocating opportunity thus involves a delicate balance between immediate and future utility.

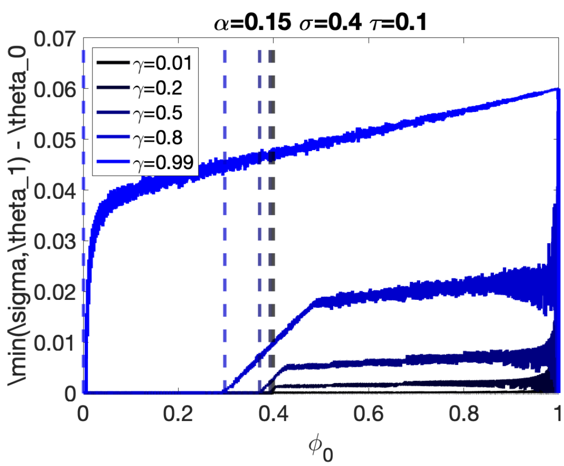

Figure 1: The difference between θ_0 and θ_1 at every state $0 \leq \phi_0 \leq 1$. The dashed lines specify the tipping points beyond which the optimal policy does not use affirmative action. Note that the extent of affirmative action is increasing in ϕ_0.

Whether socioeconomic affirmative action is employed by the optimal solution — and the extent to which it is employed — depends on the fraction ϕ_0 of individuals from group D; in the most basic model, the amount of affirmative action decreases monotonically as ϕ_0 is reduced. The extent of affirmative action is also determined by the amount of opportunity available (α), the dependence of success on ability and socioeconomic status (σ and τ), and society's patience in trading off immediate payoff in return for payoff from future generations (γ). We characterize the optimal solution in this respect as a function of these parameters, finding that for some regions of the parameter space, the society employs *temporary affirmative action*, reducing the size of group D to a given level before equalizing thresholds in subsequent generations; in other parts of the parameter space, the society employs *persistent affirmative action*, in which the threshold for group D is strictly lower in every generation and the size of group D converges to 0 over time. See Figure 1 for an illustration of the extent of affirmative action for a specif setting of σ, τ, and α.

Figure 2 provides some ways of describing the regions of parameter space in which the optimal solution uses persistent affirmative action. As the partitions of the space there make apparent, the interactions among the key parameters is fairly subtle. First, persistent affirmative action is promoted by large values of α and small values of τ, since these make it easier to include high-performing members of group D without a large difference in thresholds; and it is promoted by larger values of γ, indicating greater concern for the payoffs in future generations. One might have suspected that persistent affirmative action would only be realized in the optimal solution in the limit as society's patience (essentially $\gamma/(1-\gamma)$) goes to infinity; but in fact, a sufficiently large finite amount of patience is sufficient for the optimal policy to use persistent affirmative action.

In our model, we include a probabilistic background process by which individuals can also move between groups A and D; this reflects the idea that there are many mechanisms operating simultaneously for socioeconomic mobility, and we are studying only one of these mechanisms via the opportunity under consideration. The most basic version posits a single probability p that each individual independently loses their group membership and re-samples it from the current distribution of group sizes. We also consider a version of the model in which this probability of loss of group membership is different for groups A and D; in this case, we are only able to solve the model computationally, and these computational results reveal interesting non-monotonicities in the amount of affirmative action employed as a function of the relative size of group D (ϕ_0).

4 Concluding Discussion

Utilitarianism, Prioritarianism, and the Desert Principle. Our simple mathematical model allows us to represent and distinguish among several distinct worldviews toward allocation policies (see, e.g., [Arneson, 2013] for further discussion of these views): (1) a *utilitarian* view, which generally favors slightly lower-ability members of A to comparable, but slightly higher ability members of D in pursuit of maximizing social utility and productivity (recall that membership in A confers a boost in success probability); (2) a *prioritarian* view, which evaluates a policy according to its impact on the well-being of the worse-off members of society. Our model can capture the priority view through large discount factors (recall that as the society's patience increases, it effectively increases the priority assigned to the disadvantaged group members), or by adjusting the welfare function; (3) a *desert-principle* view, which advocates for allocating opportunities based on some notion of deservingness. Deservingness in this view is often defined in terms of the contributions people make to the social utility. Hence success probability in our model is arguably the closest match to individual desert. With that definition for desert, desert-based principles would allocate opportunities myopically in each generation. As our

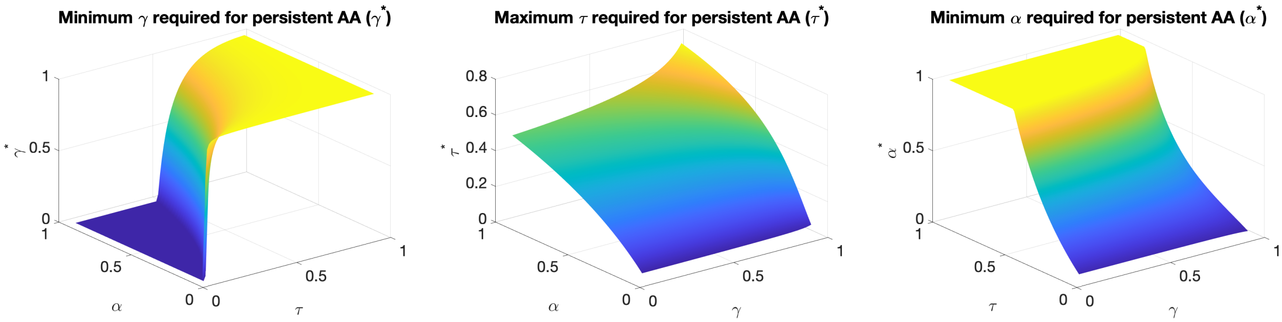

(a) (b) (c)

Figure 2: Visualizing the triples (α, τ, γ) for which the optimal policy uses persistent affirmative actions. (We set $\sigma = 1 - \tau$ for these plots.) Points that lie above the surfaces in panels (a) and (c), and below the surface in panel (b), correspond to parameter values yielding persistent affirmative action. (2a) When $1 - \frac{\alpha\sigma}{\tau} < 0$, any $\gamma > 0$ suffices for persistent affirmative action and eventually moving the entire population to group A. When $1 - \frac{\alpha\sigma}{\tau} > 0$, there exists some $\gamma < 1$ (and hence a finite level of patience $\left(\frac{\gamma}{1-\gamma}\right)$) that suffices for persistent affirmative action. (2b) When τ is sufficiently large, the optimal policy does not use persistent affirmative action; this is because for a large τ, the extent of affirmative action required to pick up the best performing members of D is large—which in turn significantly reduces the immediate payoff. For any given value of α, there exists a sufficiently small τ that guarantees persistent affirmative action. (2c) When α is small relative to τ, the optimal policy does not use persistent affirmative action; this is because the cost of picking the best performing members of D is very high and a small A group suffices for filling the available opportunities. Note that for some values of τ, no matter how large α is, the optimal policy never employs persistent affirmative action.

analysis illustrates, such policies often fail to maximize the social utility in the long-run.

Limitations and Interpretations. Our model is designed to incorporate the basic points we just mentioned in as simplified a fashion as possible; as such, it is important to note some of its key limitations. First, it is intended to model the effect of a single opportunity, and it treats other forms of mobility probabilistically in the background. It also assumes that the fundamental parameters $(\alpha, \sigma, \tau, \gamma)$ are constant over all generations as well as over individuals within one generation. It treats an individual's group membership (A and D) and ability as a complete description of their performance, rather than including any dependence on the group membership of the individual's parent. (That is, an individual in group A performs the same in the model regardless of whether their parent belonged to group A or D.) All of these would be interesting restrictions to relax in an extension of the model. Second, much of the past theoretical work on intergenerational mobility focuses on an issue that we do not consider here: the strategic considerations faced by parents as they decide how much to consume in the present generation and how much to pass on to their children. Our interest instead has been in the optimization problem faced by a social planner in allocating opportunities, treating the behavior of the agents as fixed and simple. Here too, it would be interesting to explore models that address these issues in combination. Finally, because our focus is on intergenerational mobility in a socioeconomic sense, we do not model discrimination based on race, ethnicity, or gender, and the role of race- or gender-based affirmative action in combatting these effects. The model is instead concerned with *socio-economic* or *class-based* [Malamud, 1995, Kahlenberg, 1996] affirmative action. That said, the ingredients here could be combined with models of statistical or taste-based discrimination on these attributes to better understand their interaction.

The simplicity of our model, however, does allow us to make a correspondingly fundamental point: that even a purely payoff-maximizing society can discover affirmative action policies from first principles as it seeks to optimize the allocation of opportunities over multiple generations. Moreover, the optimal allocation policy is deeply connected to dynamic programming over the generations; the society is essentially attempting to "steer" the balance of group A and group D over time, making sure not to turn things too abruptly (giving up present benefit) or too gradually (giving up future benefit). This idea that society is searching for a way to turn optimally toward a better outcome is not specific to our model; it is an image that has arisen in qualitative discourse over several centuries. It can be seen in a quote popularized by Martin Luther King, that "the arc of the moral universe is long, but it bends toward justice" [Cohen, 2006]. Interestingly, the original form of this quote, by the American minister Theodore Parker in 1853, has an even more abstractly mathematical flavor: "I do not pretend to understand the moral universe; the arc is a long one, my eye reaches but little ways. I cannot calculate the curve and complete the figure by the experience of sight; I can divine it by conscience. And from what I see I am sure it bends towards justice" [Parker, 1853]. It is a curiously apt image for the way in which our optimal solutions gradually turn through the state space to reshape the distribution of socioeconomic groups, and it can be seen as added motivation for the issues at the heart of the model.

References

[Arneson, 2013] Richard Arneson. Egalitarianism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2013 edition, 2013.

[Becker and Tomes, 1986] Gary S. Becker and Nigel Tomes. Human capital and the rise and fall of families. *Journal of labor economics*, 4(3, Part 2):S1–S39, 1986.

[Chetty *et al.*, 2014] Raj Chetty, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *American Economic Review*, 104(5):141–147, 2014.

[Cohen, 2006] Joshua Cohen. The arc of the moral universe. *Philosophy and Public Affairs*, pages 91–134, 2006.

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

[Forde-Mazrui, 2004] Kim Forde-Mazrui. Taking conservatives seriously: A moral justification for affirmative action and reparations. *California Law Review*, 92:683, 2004.

[Kahlenberg, 1996] Richard D. Kahlenberg. Class-based affirmative action. *California Law Review*, 84:1037, 1996.

[Lee and Solon, 2009] Chul-In Lee and Gary Solon. Trends in intergenerational income mobility. *The Review of Economics and Statistics*, 91(4):766–772, 2009.

[Loury, 1981] Glenn C. Loury. Intergenerational transfers and the distribution of earnings. *Econometrica: Journal of the Econometric Society*, pages 843–867, 1981.

[Malamud, 1995] Deborah C. Malamud. Class-based affirmative action: Lessons and caveats. *Texas Law Review*, 74:1847, 1995.

[Page, 2008] Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, 2008.

[Parker, 1853] Theodore Parker. *Ten Sermons of Religion*. Charles Francis and Company, 1853.