

Algorithmic Pluralism

A Structural Approach To Equal Opportunity

SHOMIK JAIN¹, VINITH M. SURIYAKUMAR², KATHLEEN A. CREEL³, and ASHIA C. WILSON^{1,2},

¹Institute for Data, Systems, and Society, MIT ²Department of Electrical Engineering and Computer Science, MIT

³Department of Philosophy and Religion and Khoury College of Computer Sciences, Northeastern University

We present a structural approach toward achieving equal opportunity in systems of algorithmic decision-making called *algorithmic pluralism*. Algorithmic pluralism describes a state of affairs in which no set of algorithms severely limits access to opportunity, allowing individuals the freedom to pursue a diverse range of life paths. To argue for algorithmic pluralism, we adopt Joseph Fishkin’s theory of bottlenecks, which focuses on the structure of decision-points that determine how opportunities are allocated. The theory contends that each decision-point or “bottleneck” limits access to opportunities with some degree of severity and legitimacy. We adopt Fishkin’s structural viewpoint and use it to reframe existing systemic concerns about equal opportunity in algorithmic decision-making, such as patterned inequality and algorithmic monoculture. By arguing for the urgent priority of alleviating severe bottlenecks, we contend that promoting algorithmic pluralism will help promote equal opportunity in a systemic way. We further show how this framework has several implications for system design and regulation through current debates about equal opportunity in algorithmic hiring.

Additional Key Words and Phrases: algorithmic fairness, bottlenecks, equal opportunity, discrimination, algorithmic monoculture, homogenization, structural injustice

1 INTRODUCTION

Almost all principles of fair decision-making derive their moral force from the principle of equality: the idea that all persons have equal moral worth and therefore should be treated equally. Since the general principle of equality could give rise to many specific types of equality, theorists of equality often debate what in particular should be equalized, whether it be opportunity, welfare, esteem, material goods, or capabilities [2, 41, 45]. This debate about what to equalize extends to the algorithmic fairness literature, with proponents of demographic parity arguing for equalizing acceptance rates [8, 51], proponents of equalized odds advocating for equalizing the likelihood of particular decision outcomes [24], and proponents of calibration arguing for equalizing the meaning of model predictions [25, 39]. Amidst this disagreement, *equality of opportunity*, the idea that everyone should have the same opportunity to succeed, enjoys a broad consensus. Equal opportunity posits a world in which circumstances of birth and parentage do not fully determine one’s life outcomes. It is a value that is singled out in several national founding documents and is regularly invoked by advocates of radically different political and social agendas [18]. The principle of equal opportunity also underlies social programs with widespread support around the world, such as public education, disaster relief, and food assistance.

What does equality of opportunity mean in the context of algorithmic decision-making? Formal views of equal opportunity often contend that all decision subjects should have a chance at a positive outcome and that their likelihood of receiving that outcome accords with their merit or desert. The fairness metric of equalized odds, for instance, checks if the predictor and protected attributes (e.g. race or gender) are independent conditional on the outcome [24]. However, a growing number of scholars [3, 20, 27, 29, 44] have criticized formal algorithmic fairness interventions that focus only on properties of model predictions and evaluate each decision-instance separately. In Ben Green’s terminology, they instead advocate for more *structural views of equal opportunity* that take into account repeated encounters with decision-making systems as well as broad patterns of life chances and pathways [20]. For example, a structural critique

put forth by Benjamin Eidelson [16] questions systems of decision-making that consistently penalize those with already poor outcomes and reward those with good outcomes. Such a system, he argues, tends to create a “patterned inequality” in which existing disadvantage causes further disadvantage, compromising the possibility of equal opportunity and leading to a variety of social-ills and overall suffering. Using algorithms to support decision-making is likely to reinforce patterns of inequality due to their ability to recognize and precisely exploit existing social patterns. Another related concern about how opportunities are structured critiques a state of affairs in which the same individuals or groups of people are on the losing side of a large number of social allocations – a phenomenon which Kathleen Creel and Deborah Hellman call “systemic exclusion” [12], Rishi Bommasani et al. call “outcome homogenization” [6] and Ifeoma Ajunwa calls “(algorithmic) blackballing” [1]. Data-driven decision-making could exacerbate these concerns due to a growing “algorithmic monoculture” [32] in which decision-makers that collectively dominate a sector rely on similar datasets and/or models and subsequently arrive at the same decisions.

In this work, we respond to the threat that patterned inequality and algorithmic monoculture pose to equal opportunity by arguing for the value of *algorithmic pluralism*. Algorithmic pluralism is the state of affairs in which the decision-making algorithms that structure opportunity are meaningfully pluralistic, which is to say that their decisions result in a plurality of paths to different outcomes, in some cases because the decisions are made on the basis of different features and values. To argue for algorithmic pluralism, we adopt Joseph Fishkin’s structural approach towards equal opportunity [17] and extend it to algorithmic decision-making. Fishkin highlights the importance of *bottlenecks*, which are decision points or “narrow places” in the structure of opportunity. Bottlenecks are the mechanisms through which opportunities are created, distributed, or destroyed. In these real-world networks of decisions, bottlenecks often chain together so that the output of one decision point is an input to the next. Each bottleneck can be evaluated in terms of its “severity,” or degree to which it constrains opportunities, and its “legitimacy,” or its justification in relation to how it allocates opportunities. Fishkin highlights that systems of decision-making with severe bottlenecks should often be considered an infringement of equal opportunity, regardless of their legitimacy. Subsequently, alleviating severity by multiplying the number of pathways to opportunity should be an urgent priority. We extend this idea of *opportunity pluralism* [17] to systems of algorithmic decision-making, resulting in our argument for algorithmic pluralism. Specifically, we make the following contributions:

- We review classic tensions in the equality of opportunity literature to motivate Fishkin’s idea of opportunity pluralism (Section 2).
- We present Fishkin’s theory of bottlenecks and use the framework to center several structural concerns about equal opportunity raised by a variety of scholars who focus on algorithms and issues of justice (Section 3).
- We propose algorithmic pluralism, highlighting the importance of alleviating severe bottlenecks in systems of algorithmic decision-making (Section 4).
- We motivate and illustrate this framework through current debates about equal opportunity in algorithmic hiring (Section 5).
- We end our discussion with implications for regulation and design (Section 6).

2 EQUAL OPPORTUNITY AND ITS LIMITATIONS

In order to motivate Fishkin’s idea of opportunity pluralism, we will briefly explore classic tensions in the equality of opportunity literature that highlight the importance of considering both access to opportunity and the structure of opportunity itself. We begin by sketching a classic thought experiment proposed by Bernard Williams [50] and discussed

by Fishkin [17]. The thought experiment asks us to imagine a society in which there are only two social roles: warriors and non-warriors. In this society, warriors are highly valued. They receive all the material and non-material goods the society has, such as superior lodging, food, salary, opportunity for advancement, and social esteem. Since there are a fixed number of positions in the army, not all members of the society can become warriors. A meritocratic process is established to determine who can become a warrior: at the age of sixteen, all members of the society participate in a multi-day competition that is intended to measure characteristics such as athletic ability, endurance, bravery, and battlefield strategy.

Imagine that an algorithm is created that uses the results of the competition to determine who should become a warrior. Under one interpretation of equal opportunity, a fair algorithm is one that gives all members of the society an equal opportunity to become a warrior by choosing warriors based only on features relevant to their success on the battlefield. However, is this algorithm truly fair? After all, children of warrior parents have a significant advantage. Since their parents are warriors themselves, they are better able to train their children in the martial arts. And because their parents are the wealthiest and most esteemed members of society, their children are more likely to be better fed, healthier, stronger, and more confident – all traits that are reflected in the differential success rate of children from warrior and non-warrior parents. The children of warrior parents also benefit from additional experiences and resources, such as family wealth and access to better education, that better prepare them for the warrior test.

These *developmental opportunities* make it impossible for warrior and non-warrior children to have a fair contest. While non-warrior children are not explicitly excluded because of their caste, the data collected on them reflects their lack of developmental opportunities. For this reason, selecting warriors based on who will perform best without equalizing their childhood opportunities in any way will tend to disadvantage non-warrior children. As Fishkin explains, “if success breeds success, and we reinforce achievement with new and richer developmental opportunities, then the project of equalizing opportunities comes squarely into conflict with rewarding performance. In that case, the very earliest developmental opportunities, which precede any meaningful performances worth rewarding, begin to take on an outsized significance” [17].

The problem of a starting gate.

Given that childhood disadvantage becomes disadvantageous [16], making later meritocratic contests unfair, one of the most important principles of equal opportunity has been to equalize developmental opportunities by providing all children (and later, all people) with equal access to education, nutrition, and other building blocks of development. This approach is sometimes called equalizing the *starting gate* in the race of life, ensuring that all the competitors show up with equal opportunity to prepare with the hopes that decision-making on the basis of merit after that point will be fair as a result [41]. Since all agree that it is impossible to fully equalize the starting gate – even with equal material opportunity, some parents are kinder and more supportive than others – some scholars have even attempted to re-create an equal starting gate through statistical methods [23].

The problem of isolating merit.

Allocating goods on the basis of later merit without equalizing developmental opportunities undermines the normative rationale for allocating goods on the basis of merit at all. If one justification for allocating goods on the basis of merit is that it rewards hard work and dedication, a profoundly unequal playing field makes it impossible to isolate those qualities from advantageous circumstances of birth and life experiences. As Fishkin notes, “everything we are and everything we do is the product of layer upon layer of interaction between person and environment – between our selves, our efforts, and our opportunities – that in a sedimentary way, over time, build each of us into the person we become” [17]. Because

of this, any real-world data and features collected to become the basis of algorithmic decision-making will inherently combine aspects of hard work, innate talent, and developmental opportunities in ways that are difficult to distinguish.

The problem of metrics and what is measured.

Another challenge to the fair implementation of equal opportunity is that any test established to measure merit is unlikely to perfectly measure real-world performance. Warrior children may score higher on any *test* of athletic ability, for example, because of they have received special coaching that improves their scores on the test more. However, this coaching may improve test scores more than it improves real-world battlefield performance. In general, given it is impossible to measure merit precisely, proxy metrics will be used which will inevitably unfairly limit the opportunity of some.

The problem of zero-sum thinking.

Focusing on specific decision points, such as the test of sixteen year-old warriors, also needlessly reduces considerations of equal opportunity to zero-sum thinking. For example, if we choose to give compensatory bonus points to non-warrior children, some warrior children may object at no longer being selected. This distributive view forces zero-sum trade-offs between mathematically incompatible notions of fairness, magnifying the stakes of choosing between these definitions [31]. As Green argues, restricting analysis to specific decision points also cannot fully “account for the inequalities that often surround those decision points [and] is therefore prone to reproducing existing patterns of injustice” [20]. By expanding our view to the broader structure of opportunities, we can better account for the important contexts in which decisions are made and consider interventions that diminish or even eliminate the zero-sum framing.

The warrior society thought experiment illustrates that focusing only on fairness metrics takes an overly narrow view of the opportunity structure. As Dewey astutely points out, “the way in which [a] problem is conceived decides what specific suggestions are entertained and which are dismissed” [15]. Formal fairness methods often reduce considerations to relative advantages or disadvantages at a particular decision point. This type of single-axis thinking comes at the expense of attention to what produces the systematic benefits or privileges in the first place [27]. In doing so, formal approaches do not account for the fact that having a plurality of opportunities (and not just equal opportunities) is a prerequisite for the accompanying virtues of freedom and the ability to shape our lives that make equal opportunity an attractive value in the first place.

3 BOTTLENECK THEORY: A STRUCTURAL VIEW OF EQUAL OPPORTUNITY

So far we have seen that establishing a fair starting gate and fair metrics for performance are *necessary* for equal opportunity in a strict meritocratic society. Next, we follow Fishkin in arguing that they are not *sufficient*.

3.1 Opportunity Pluralism

The warrior society thought experiment reveals that equalizing developmental opportunity alone is insufficient for maintaining the virtues associated with equal opportunity. Imagine we can fully equalize developmental opportunities by placing all children into warrior skill academies from the moment of birth. While those not selected to become warriors now have equal developmental opportunities, they still may have a broader complaint about the *structure of opportunity* in their society: that the bottleneck through which they must pass to succeed is too narrow. In Fishkin’s terms, a society whose pathway to success is *only* open to those with the same set of natural talents (e.g. javelin throwing, martial courage, and battlefield tactics) has established an overly “severe bottleneck” on opportunity [17].

Such a severe narrowing of life paths compromises the notion that equal opportunity is a principle which offers the freedom to shape our lives. That is, it is not enough that opportunities be equal. For example, in the warrior society, children with intellectual capacities should have the opportunity to become scribes or advisors and children with creative abilities and dexterity should have the opportunity to become craftspeople. Given that natural talents and inclinations are not chosen, akin to the lack of choice in one’s parents, a society that provides equal opportunity should provide many ways in which natural talents can lead to success. There must be a plurality of opportunities, not merely equal access to one opportunity. Even more broadly, there must be multiple ways of assessing capacities for opportunities. Having only one test such as in the warrior society establishes a severe bottleneck centering on the method of assessment. Its flaws, weaknesses, and oversights become magnified, meaning that even a child who could be a talented warrior might have their skills overlooked if they are not precisely the skills measured by the test. In addition to a plurality of opportunities, there must be a plurality of assessments. These considerations lead Fishkin to propose four central principles of *opportunity pluralism* [17]:

- (1) There should be plurality of values and goals.
- (2) As many as possible of the valued goods should be less positional and the valued roles less competitive.
- (3) As far as possible, there should be a plurality of paths leading to these different valued goods and roles, without bottlenecks constraining people’s ability to pursue those paths.
- (4) There should be a plurality of sources of authority regarding the elements described in the other principles.

3.2 Bottlenecks

Fishkin’s opportunity pluralism [17] centers on alleviating bottlenecks, or “narrow place[s] in the opportunity structure”. A bottleneck is a crucial decision-point that affects how future opportunities are created, distributed, and controlled [17]. Examples of bottlenecks include decision-points that allow one to gain credentials or develop skills that are required to pursue certain opportunities. Bottlenecks are often chained together given that in systems of decision-making, the outcomes of upstream decisions usually factor into downstream decisions. For instance, the college admissions process constitutes a significant bottleneck for many employment opportunities and life paths, but is also in itself controlled by the bottlenecks of standardized testing and access to extra-curricular activities. As algorithms are increasingly used to inform or make decisions about our lives, they can become bottlenecks that control access to many real-world opportunities such as employment, housing, and education. Legibility to resume screening algorithms is a bottleneck on the path to employment; approval by tenant screening algorithms is a bottleneck on the path to housing; and acceptance by college admissions algorithms may soon be a bottleneck on the path to certain elite institutions. In his framework for opportunity pluralism, Fishkin proposes the qualitative measures of *severity* and *legitimacy* to assess the impact of each bottleneck on the broader opportunity structure. We show how these measures map to several predominant concerns about algorithmic decision-making and its potential to infringe upon structural access to opportunity.

3.3 Bottleneck Severity

The severity of a bottleneck is the degree to which it constrains opportunities. Severity is a combination of two factors: pervasiveness and strictness [17]. A bottleneck’s *pervasiveness* is the range of people and opportunities it affects. For example, the warrior selection algorithm in the warrior society has maximum pervasiveness because it controls everyone’s access to all the valued goods in the society. *Strictness* refers to how directly a bottleneck controls opportunities, or in other words, whether it is an absolute bar, strong preference, or weak preference. The warrior




Fig. 1. (Fig. 6 in [17]) Fishkin’s framework for classifying bottlenecks. He proposes mapping the structural impact of bottlenecks along the axes of severity and legitimacy.

selection algorithm also has maximum strictness because it is an absolute bar: no one can become a warrior unless they are selected by the algorithm. While the warrior society represents an extreme case, many severe bottlenecks naturally arise in systems of data-driven decision-making, as illustrated by the concerns expressed in the ideas of *patterned inequality* and *algorithmic monoculture*.

The concern about patterned inequality is a concern about severe bottlenecks.

Patterned inequality refers to two observations made by Benjamin Eidelson: (1) real-world inequalities in status, resources, and opportunities are often patterned in terms of certain discernible, socially salient features; (2) allocating future outcomes and resources based on an individual’s likelihood of success will predictably reproduce and aggravate patterns of this kind [16]. The problem of developmental opportunities makes it inevitable that some patterns of inequality will exist. Once an observable trait, perhaps implicit or arbitrary in itself, is correlated with less visible attributes relevant for deciding opportunities, decision-makers will rationally tend to use that trait as a proxy for allocating future outcomes (c.f. the problem of isolating merit). For example, consider the strong correlation between credit score and job performance, which led to its controversial use as a proxy for hiring decisions [4]. Given algorithms are developed to identify these patterns in a more explicit way, the concern is that allocating opportunities using them “will tend to reproduce existing patterns in inequality and cement the matrix of stereotypes and social meanings that both cause and result from those patterns” [16]. This stronger tendency towards the comparatively privileged creates the possibility of stricter bottlenecks, while the denial of even more opportunities to members of worse-off groups represents more pervasive bottlenecks.

The concern about algorithmic monoculture is a concern about severe bottlenecks.

Algorithmic monoculture [32] occurs when multiple decision-makers controlling access to a large quantity of valued goods rely on the same or similar datasets and/or models. As Bommasani et al. [6] discuss, algorithmic monoculture can lead to a pattern of homogenous outcomes in which the same people are subject to consistent errors or negative outcomes. Creel and Hellman [12] have argued that at scale, this “systemic exclusion” is unfair to decision-subjects even if overall accuracy is high. Algorithmic monocultures are increasingly common in high-stakes screening decisions across many domains, including in employment, healthcare, lending, and criminal justice [38]. One reason why monoculture occurs is that few decision-makers have the resources to build their own automated systems. For example, over 30% of

Fortune 100 companies use the same resume-screening service for hiring [40]. Recent works have also suggested that the performance gains of using pre-trained foundation models could lead to many decision-makers relying on the same or similar foundation models, leading to standardized errors [7, 46]. In other domains, the data collection itself may be expensive or unfeasible, such as in healthcare algorithms which disproportionately use data from a handful of large hospitals in only three states: California, Massachusetts, and New York [12].

The standardization of decision-making processes could create stricter bottlenecks because certain features will be stronger preferences for the opportunities at hand. In ranking problems like comparing job candidates, sharing an algorithm across multiple decision-makers can also reduce the overall quality of decisions, even if the algorithm is more accurate for any one decision-maker in isolation [32]. Moreover, if similar algorithms are uniformly applied across wide swathes of a single domain, outcomes could become homogenized resulting the systematic exclusion of groups and individuals from opportunities of great consequence [6, 12]. Thus outcome homogenization constitutes a pervasive bottleneck for those that repeatedly receive undesirable decisions because they may find themselves locked out of many paths in the opportunity structure.

3.4 Bottleneck Legitimacy

Other structural concerns about equal opportunity center on a model’s *legitimacy*, or its justification in relation to how opportunities are allocated [17]. The justification for an algorithm is often argued on two grounds: mathematical legitimacy and social legitimacy. *Mathematical legitimacy* is simply the algorithm’s accuracy with respect to a particular problem formulation. As we saw in the warrior society, however, a high accuracy alone does not fully justify the use of an algorithm to control real-world outcomes. The other necessary aspect to consider is an algorithm’s *social legitimacy*: what is the rationale for using a particular problem formulation (e.g. choosing a specific model, features, or training objective) to constrain the specific opportunity at hand? While many scholars have raised concerns about legitimacy of algorithmic decision-making procedures [21, 27, 29], we highlight two in particular that direct their legitimacy critique at the way opportunities are structured.

The concern about compounding injustice is a concern about illegitimate bottlenecks.

Compounding injustice refers to Deborah Hellman’s critique of decisions that take a prior wrong that a person has suffered, or its effects, as a reason for allocating future opportunities in a way that makes them still worse off [26]. In many cases, adverse predictions by algorithms may involve the use of features that relate to prior victimization. For example, a prominent healthcare algorithm selecting patients for high-risk care programs used prior health costs as a proxy for future health needs [37]. This ignored the well-documented injustices of the U.S. healthcare system that cause less money to be spent on black patients who have the same level of need as equally sick white patients. By denying future care on the basis of prior health costs, the algorithm “compounded the initial injustice” of unequal access to care. From the perspective of bottleneck theory, Hellman’s formulation of compounding injustice and Eidelson’s patterned inequality both surface similar concerns about the structure of opportunities; however, Hellman’s argument raises the illegitimacy of the initial condition as central to moral considerations, whereas Eidelson centers the severity of unequal patterns of opportunity as sufficient to warrant moral concern.

The concern about relational harms is a concern about illegitimate bottlenecks.

Relational harms refer to when algorithms fail to account for the power relations, social dynamics, and structural contexts that surround systems of decision-making in the real-world [20, 21]. In particular, this involves scrutinizing what institutions, values, and norms cause social and material disparities in opportunities [2, 35]. Explicitly considering

the potential for relational harms can help justify algorithmic design choices and problem formulations. Whether or not certain input features and output predictions are legitimate, for instance, can depend on if they reduce dignitary and material disparities that reflect social hierarchies [20]. For example, arguments against the legitimacy of COMPAS highlight that predictions of recidivism fail to account for the racial biases in policing that make re-arrest rates higher for minority populations [10]. In addition, how much a decision-maker values mitigating various relational harms can also inform the choice of a training objective [14]. Companies making hiring decisions may have different values such as hiring from the local community, creating more diverse teams, or addressing past discrimination. Training an algorithm based on who has been hired in the past may undervalue the contributions that underrepresented applicants bring to a company. Ultimately, considerations of these broader contexts and how they affect an algorithm’s social legitimacy will depend on one’s choice of worldview [19]. This may differ from the perspective of various stakeholders, such as those making decisions versus those seeking opportunities.

4 ALGORITHMIC PLURALISM

In this section, we draw on the concept of opportunity pluralism to define *algorithmic pluralism*. We first argue that alleviating severity is the most important problem to address in order to make opportunities more equal. The previous section showed that Fishkin’s structural view of equal opportunity requires making bottlenecks both less severe and more legitimate. But as many including Fishkin have argued, severity alone can be sufficient to warrant intervention even if a bottleneck is legitimate. We then define algorithmic pluralism by considering different ways in which an algorithm could be pluralistic. Ultimately, we conclude that of these different kinds of pluralism, the most important is that algorithms allow individuals with diverse characteristics (skills, identities, life paths, etc.) to have a variety of chances to gain access to valued opportunities and goods.

4.1 Legitimacy Matters, But Severity Takes Precedent

The importance of ensuring legitimacy in algorithmic decisions cannot be overstated. Legitimacy aims to ensure that algorithmic decision-making is acceptable or justified. An algorithm can achieve legitimacy by being produced in accordance with legitimate process, such as a democratic procedure, or by its alignment with fundamental democratic values such as accuracy, equality, fairness, and justice [11, 28, 30]. Interventions to increase algorithmic legitimacy that rely on achieving “fairness” have been a central focus of the algorithmic fairness community. Recent works rightfully encourage the algorithmic fairness community to consider the legitimacy of the broader opportunity structure in addition to the fairness of the algorithm itself [20, 27, 29]. The convergence of this work suggests that an algorithm that lacks mathematical or social legitimacy is in need of intervention.¹ If an algorithm or its decision outcomes are not legitimate, then those subject to the algorithm have the right to demand that it become legitimate, or that it cease to have the power to affect their lives. Thus legitimacy alone can be a sufficient condition for intervention.

We can determine whether an individual decision or a decision-making system is legitimate without evaluating the broader structure of opportunity in which the decision takes place. Severity, by contrast, can be observed solely by evaluating the influence of decision-makers on the ultimate pattern of outcomes for decision subjects. This leads Fishkin to center severity as the prioritized axis for intervention in order to bring about opportunity pluralism. By intervening

¹As Ben Green points out, there exists significant political disagreement about what makes a decision legitimate. Resolving the disagreement involves “grappling with contested notions of what types of inequalities are unjust and what evidence constitutes sufficient proof of social hierarchy” [20, p90]. Although we consider this grappling to be important and necessary work, for the purpose of this paper we will rely on Fishkin’s “mathematical legitimacy” and “social legitimacy”, introduced in the previous section. However, we believe that our arguments would hold were other notions of legitimacy to be substituted. For discussions of algorithmic legitimacy, see [28, 48, 49].

to make existing bottlenecks less severe, it becomes “more possible, both in ideal theory and the real world, to achieve the goals that lead us to value equal opportunity in the first place” [17]. In algorithmic systems, we similarly contend that severity warrants intervention on its own, as other works have argued using different frameworks [6, 12, 16].² For example, an algorithmic monoculture can cause severe restrictions on the space of opportunity [32]. If the same algorithm, or a set of very similar algorithms, makes decisions across a region, they can create severe bottlenecks that homogenize outcomes [6]. Regardless of how legitimate these outcomes are considered to be, the consistency of algorithmic exclusion due to algorithmic monoculture is a moral concern, as Creel and Hellman argue [12]. In the next section, we extend the idea that alleviating severity should be an urgent priority to systems of algorithmic decision-making and show that alleviating severity requires *algorithmic pluralism*.

4.2 Pluralism in Algorithmic Decision-Making

Algorithmic pluralism describes a state of affairs in which the algorithms used for decision-making are not so consistent as to constitute a severe bottleneck on opportunity. In defining algorithmic pluralism, we extend opportunity pluralism to systems of algorithmic decision-making. However, achieving algorithmic pluralism requires resolving an important question not fully addressed in Fishkin’s text: what exactly must be plural about algorithmic decision-making? Elaborating on the principles of opportunity pluralism described in subsection 3.1, pluralism in algorithmic systems could entail:

- (1) pluralism of algorithmic components (e.g. models, features, and training objectives);
- (2) pluralism of evaluation criteria used to determine who receives each opportunity [17, p16];
- (3) pluralism of algorithmic decision-making processes;
- (4) pluralism of algorithmic decision-makers [17, p16];
- (5) pluralism of opportunity: algorithms allow individuals with diverse characteristics, such as skills, identities, and life paths, to have a variety of chances to gain access to valued opportunities and goods.

In what follows, we will argue that (5) is a necessary condition for algorithmic pluralism and that while all five forms of pluralism can be valuable, (1-4) are valuable insofar as they bring about (5).

The most natural way to bring about a pluralistic opportunity structure in our society would be to encourage the proliferation of diverse institutions and decision-makers, each of which grants access to a certain kind of opportunity. If the diverse institutions and decision-makers each make their decisions on different bases – perhaps one evaluates numerical ability and teamwork, while another evaluates geometric and spatial reasoning and manual dexterity – then the opportunity structure permits many different people to reach to many different valuable life paths. Therefore, all five kinds of pluralism initially appear to be linked: access to many different choice-worthy life paths³ for (5) many different kinds of people are granted by (4) diverse institutions and (3) decision-making processes that evaluate those people on the basis of (2) many different criteria. For example, (5) aspiring warriors, musicians, and accountants, each with different skills and characteristics, might find opportunities at (4) the military, the orchestra, and the tax-preparation

²As Fishkin notes, an algorithm that is severe and legitimate is deserving of more scrutiny and moral concern because of its weight on the opportunity structure than one that is illegitimate but not severe. Importantly, many works advocating substantive algorithmic fairness (e.g. Green [20]) fail to make a distinction between prioritizing legitimacy and severity concerns.

³Whether meaningfully different life paths exist affects whether different people can flourish in the opportunity structure. Furthermore, the structure of opportunity can encourage people to reflect on and develop their own values and aspirations, or conversely can suppress such exploration by ensuring that only a limited number of life paths are rewarded with the basic conditions of human flourishing[17, p17], as in the warrior society. These two forms of pluralism are important in Fishkin’s account. However, it is not clear that algorithmic decision-making plays any special role in bringing about these features of the opportunity structure. The structure of opportunity at the society level, including which social roles exist and are rewarded or dis-valued, is rarely determined by the choice of algorithms or the choice between using algorithmic and human decision-making. Therefore we think that pluralism about values and goals are not relevant to determining the content of algorithmic pluralism.

firm on the basis of their (2) fighting prowess, musical ability, and numerical precision, respectively. It is therefore not surprising that Fishkin consistently links these different types of pluralism in his framework.

However, in algorithmic decision-making, these types of pluralism need not be linked. And if they are uncoupled, it is not clear what algorithmic pluralism requires. To make this vivid, imagine a society in which the pluralisms are not coupled: there is only one algorithm through which all must access opportunity, but since the society has vast abundance and distributes it almost equally, the top 99% of people according to the algorithmic scores are rewarded with good material life outcomes and the ability to choose a meaningful job or life pursuit. However, the algorithm is not static: it changes each year such that a different 99% are chosen (and therefore a different 1% are excluded). There is only one decision-maker and the algorithm applies a single evaluation criterion to everyone, so the society lacks pluralism of types (1-4). But the society supports pluralism of type (5): many different kinds of people have multiple chances to gain opportunities they value.

To put this question another way, is it enough to ensure that a diverse group of people receive opportunities they value? Or must they also be selected for those opportunities for a plurality of reasons and/or by a plurality of decision-makers? This question matters for algorithmic decision-making. Work in algorithmic fairness has investigated whether algorithmic biases create a patterned inequality [16] by acting as “moderately strict but pervasive” bottlenecks [17, p164] (pluralism 2) that disproportionately exclude members of marginalized subgroups (pluralism 5), but this work does not typically evaluate similarity between decision makers (pluralism 4). And work on algorithmic monoculture and outcome homogenization focuses on the extent to which similarity between decision-makers (pluralism 4) results in the same individuals being rewarded or excluded (pluralism 5) [6, 12, 33], but it does not typically investigate whether the decisions are being made on the basis of the same criteria “under the hood” (pluralism 2). Both literatures find pluralism (5) important, but they differ in their focuses on pluralisms (2) and (4).

The choice of pluralism(s) makes a significant difference to the requirements of algorithmic pluralism. If algorithmic pluralism requires that different algorithms decide based on different criteria (pluralism 2), then determining whether we are in a state of algorithmic pluralism might require access to the criteria “on the basis of which” each model made its decision as well as the ability to explain the relevance of each criterion. Achieving transparent access to these criteria in complex machine learning models has been notoriously challenging [13, 22, 36]. It would be necessary to not only achieve access to individual decision criteria, but also to have a systematic way to compare the criteria and to measure the diversity of decision criteria across the opportunity structure.

If algorithmic pluralism instead means ensuring that many different decision-makers control access to opportunities (pluralism 4), we would be most interested in evaluating *algorithmic monoculture* and the extent to which decision-makers within the ecosystem resemble one another [32]. If we instead required only that the opportunity structure contains multiple different decision-making organizations, each organization could rely on the same decision-making algorithm, rendering their decisions identical to one another. In order for pluralism 4 to be meaningful, the decision-makers must be different at the level of the algorithm used, not merely different in organization name.

Pluralism of type 5 requires that the opportunity structure be devoid of severe bottlenecks so that a broad set of individuals can access opportunity. The winners are not the only ones who can keep winning; those unfortunate in one contest have other chances to flourish in the future. As Creel and Hellman argue, hiring algorithms can achieve this kind of pluralism by allowing different kinds of job candidates to succeed, neither consistently making mistakes on the same individuals nor establishing only one criterion for their success [12]. Due to the consistency of algorithmic decision-making, this type of algorithmic pluralism is especially important in ensuring that many individuals have access to the opportunities that would allow them to flourish.

4.3 Algorithmic Pluralism Requires Pluralism of Opportunity

We argue that only (5) is *necessary* for algorithmic pluralism. If we ultimately hope to alleviate severe bottlenecks and promote opportunity pluralism, a range of people must actually have access to and receive different forms of opportunity that they value. In other words, algorithms must allow a pluralism of opportunity (5). Pluralisms (1-4) are *instrumentally* valuable, which is to say they are valuable only insofar as they achieve the goal of (5). Having diverse model components and evaluation criteria across different decision-makers and processes is not valuable on its own – it is only valuable if it leads to diverse model outcomes, each of which opens doors to opportunity to different people. If the diverse models all deliver opportunity to the same small group of people, the diversity of their individual components is not materially meaningful. In short, algorithmic pluralism requires algorithms to differ from one another in their decision-outcomes, often because they differ in the bases for their decisions, allowing decision-subjects multiple kinds of opportunity.

Achieving pluralism of opportunity (5) in practice might require bringing about one or more of pluralisms (1-4): algorithms different from one another (4) in their components (1) in their evaluation criteria (2), in the broader decision-making processes surrounding them that lead to outcomes (3). As we will discuss in Section 6, regulators and designers should often seek to achieve pluralism by promoting (1-4). However, it is also possible for a single algorithm to bring about (5) if (for example) it incorporates randomness in its decision-making to broaden the diversity of people who are given access to valued outcomes in a series of decision instances over time [12]. What ultimately matters is that the resulting opportunity structure is meaningfully pluralistic in offering a diverse set of individuals with access to opportunity. In the next section, we use a case study of algorithmic hiring to further illustrate the idea of algorithmic pluralism.

5 CASE STUDY: ALGORITHMIC HIRING

We elaborate on the importance of algorithmic pluralism with a case study of algorithmic hiring. Algorithms constitute severe bottlenecks across all stages of contemporary hiring. Resume-screening and skill-assessment algorithms often reject the majority of applicants, making them strict bottlenecks on the path to receiving an interview. Moreover, most Fortune 500 companies rely on hiring tools from third-party vendors such as HireVue and Pymetrics [40], creating the risk of the same algorithms being pervasive bottlenecks for multiple jobs, strictly excluding the same candidates.

Hiring algorithms have the potential to reinforce existing patterns of inequality in the employment sector. Decades of audit studies show that employers tend to discriminate against racial and gender minorities, with little improvement over the past 25 years [40]. Unsurprisingly, Amazon’s original resume-screening algorithm downgraded applicants from women’s colleges, in part because it was trained on its existing majority-male employee base [42]. LinkedIn’s job recommendation algorithms also ended up referring more men than women for open roles because it tried to maximize applicants and men were more likely to apply even if they didn’t meet the qualifications [43]. In addition to replicating the biases long demonstrated in human hiring, algorithmic hiring has the potential to establish a uniquely severe bottleneck to employment. If many employers rely on the same large third-party provider of candidate screening software, the same individual candidates or subgroups could be consistently and erroneously rejected [6]. This targeted exclusion could be both severe and illegitimate, due to the arbitrariness of the factors on which candidates are rejected [12].

Disparate impact regulations in the United States aim to loosen bottlenecks for legally protected groups, but existing regulations fall short of alleviating the bottleneck imposed by hiring algorithms. The U.S. Equal Employment

Opportunity Commission (EEOC) has established the 4/5 rule as precedence for when a disparate impact case can be brought against an employer [9]. The rule takes a formal view on equal opportunity and states that employers may be at legal risk of discrimination if the selection rate for one protected group (e.g. race, gender, disability, age, religion, etc.) is less than 4/5 of that of the group with the highest selection rate. However, employers can still maintain their selection procedures by showing a “business necessity” [9]. This often leads to zero-sum debates in the courts over what procedures are legitimate, where the status-quo benefits one group and the alternatives would distribute opportunities to another group. Ultimately, neither option offered by the current regulations fully addresses the strictness or pervasiveness of bottlenecks created by algorithmic hiring procedures or by algorithmic monoculture in hiring procedures.

The EEOC could adopt several interventions in order to promote algorithmic pluralism. As third parties with the power to audit and monitor firms, regulators are in the best position to measure the homogeneity of hiring outcomes, since individual companies cannot fully observe the decisions made by their competitors. This could entail auditing information about which candidates are rejected and the extent to which these outcomes are systemic and algorithmically driven. For example, resume-screening algorithms may reject individuals who have employment gaps in their resume, even if they had to stop working for legitimate reasons like childcare or health problems. While the 4/5 rule would protect applicants whose unemployment status was caused by a disability, audits could uncover other groups of unemployed persons that are systemically rejected. Moreover, protections for these groups need not be zero-sum: policies could mandate that employers provide ways for candidates with special circumstances to flag these in their application, which would help alleviate the strictness of algorithmic-screening decisions for all applicants.

Employers can further shift their hiring processes towards algorithmic pluralism and may even have incentives to do so. If all competitors rely on the same hiring algorithm, then the overall quality of hired applicants can decrease over time because those selected by the algorithm can only choose one company to work for [32]. Employers may also have different conceptions of applicant merit or business priorities that a standardized algorithm will fail to capture. To address these problems, companies could modify the design of hiring algorithms to prioritize certain features such as skills on a resume or add training objectives like hiring from the local community. Even if companies opt to use pervasive algorithms for efficiency or cost tradeoffs, the EEOC should consider stepping in to prohibit anti-competitive practices, especially if direct competitors are using the same hiring algorithms. Moreover, many algorithms will undervalue applicants from under-represented groups and fail to learn about changes in applicant hiring potential over time. This should incentivize employers to not strictly follow algorithmic rankings and balance exploitation (selecting from groups with proven track records) with exploration (selecting from other groups to learn about quality) [34].

6 IMPLICATIONS FOR ALGORITHMIC REGULATION AND DESIGN

Regulators and policymakers in many branches of government are explicitly tasked with equalizing the structure of opportunities. In systems of algorithmic decision-making, regulatory intervention could entail better mechanisms to identify severe bottlenecks and legal protections or recourse for those whose opportunities are most severely constrained. While individual decision-makers may not be able to fully observe the bottlenecks created by others in their ecosystem, they can still leverage model multiplicity to customize algorithms, and may even have incentive to do so.

6.1 Regulators and Policymakers

Regulators and policymakers have the unique capacity to monitor system-level outcomes and identify severe bottlenecks. For algorithmic systems in particular, this may first require the development of mechanisms to quantify severity. One

such metric could involve the extent to which deployed models make homogeneous predictions [6, 46]. However, many current systems of decision-making lack transparency about what stages of the process are algorithmically-mediated and the features of individuals that are denied opportunities. Regulators are in a unique position to collect this information through audits or other mandates while still protecting privacy and competitive advantages. Existing anti-competitive laws could also help uncover the pervasiveness of certain algorithms or models.

Policymakers can subsequently intervene in algorithmic systems that amount to a severe bottleneck, and attempt to alter those systems so that the bottleneck is somewhat less severe. For instance, many federal and state anti-discrimination laws extend varying levels of protections to groups that we would not ordinarily expect the law to show solicitude, such as those with low-credit, predispositions to disease, and even prior criminal convictions [17]. While businesses and other decision-makers may have many legitimate reasons for denying opportunities on the basis of these characteristics, policymakers consider these and other categories (as empirical matters in our society) to severely constrain a person’s range of opportunities making it sufficient and necessary to warrant concern. This protected list of categories should continue to expand over time in the advent of newly constrained groups from algorithmic decision-making. Policymakers can also proactively mandate avenues for recourse or redress [47], especially for cases in which a lack of transparency makes it difficult to identify those whose opportunities are severely constrained.

6.2 Designers and Decision-Makers

Opportunity pluralism calls for decision-makers to embed their unique values and preferences in the design of algorithmic systems. In practice, many procedural design choices exist across selecting features, training objectives, and what stages of the process are algorithmically-mediated. The same problem formulation can also yield model multiplicity [5], or the existence of multiple models with varying predictions but similar accuracy. If different decision-makers adopt a plurality of design choices, it could counteract the risks of algorithmic monoculture and alleviate the severity of decisions. Moreover, this could even increase legitimacy by improving the overall quality of decisions [32] and better accounting for long-term impacts [34]. Despite these advantages, resource constraints may still cause decision-makers to converge on similar models. For these settings in particular, Creel and Hellman propose intentionally introducing randomness into the decision-making process [12]. We encourage future work to develop more technical solutions that can help disrupt the severity of decisions under various tradeoffs.

7 CONCLUDING REMARKS

Fishkin’s idea of opportunity pluralism makes the case that societies ought to move their structures of opportunity towards a more pluralistic model, where there are many gatekeepers and paths towards opportunities [17]. As our economy increasingly relies on artificial intelligence, we emphasize the importance of extending this idea to systems of data-driven decision-making through algorithmic pluralism. Towards the end of his book, Fishkin emphasizes that opportunity pluralism has vast implications for various institutions and stakeholders. He aspires to encourage gatekeepers who wish to help build a more pluralistic opportunity structure to reexamine and ameliorate the bottlenecks that result from how they make decisions. We similarly implore the data scientists responsible for designing models, gatekeepers who ultimately make decisions, and policymakers regulating algorithmic systems to isolate where they have the leverage to ameliorate bottlenecks, especially those that are most severe.

REFERENCES

- [1] Ifeoma Ajunwa. 2021. An Auditing Imperative for Automated Hiring Systems. *Harvard Journal of Law & Technology* 34, 2 (2021).
- [2] Elizabeth S. Anderson. 1999. What Is the Point of Equality? *Ethics* 109, 2 (Jan. 1999), 287–337. <https://doi.org/10.1086/233897>
- [3] Falaah Arif Khan, Eleni Manis, and Julia Stoyanovich. 2022. Towards Substantive Conceptions of Algorithmic Fairness: Normative Guidance from Equal Opportunity Doctrines. In *Equity and Access in Algorithms, Mechanisms, and Optimization* (Arlington, VA, USA) (EAAMO '22). Association for Computing Machinery, New York, NY, USA, Article 18, 10 pages. <https://doi.org/10.1145/3551624.3555303>
- [4] Joshua Ballance, Robert Clifford, and Daniel Shoag. 2020. “No more credit score”: Employer credit check bans and signal substitution. *Labour Economics* 63 (2020), 101769.
- [5] Emily Black, Manish Raghavan, and Solon Barocas. 2022. Model Multiplicity: Opportunities, Concerns, and Solutions. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 850–863. <https://doi.org/10.1145/3531146.3533149>
- [6] Rishi Bommasani, Kathleen A. Creel, Ananya Kumar, Dan Jurafsky, and Percy S Liang. 2022. Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 3663–3678. https://proceedings.neurips.cc/paper_files/paper/2022/file/17a234c91f746d9625a75cf8a8731ee2-Paper-Conference.pdf
- [7] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [8] Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. 2009. Building classifiers with independency constraints. In *2009 IEEE international conference on data mining workshops*. IEEE, 13–18.
- [9] US Equal Employment Opportunity Commission. 1979. Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures. <https://www.eeoc.gov/laws/guidance/questions-and-answers-clarify-and-provide-common-interpretation-uniform-guidelines>
- [10] Sam Corbett-Davies, Emma Pierson, Avi Feller, and Sharad Goel. 2016. A computer program used for bail and sentencing decisions was labeled biased against blacks. It’s actually not that clear. *Washington Post* 17 (2016).
- [11] Sasha Costanza-Chock. 2018. Design justice, AI, and escape from the matrix of domination. *Journal of Design and Science* 3, 5 (2018).
- [12] Kathleen Creel and Deborah Hellman. 2022. The Algorithmic Leviathan: Arbitrariness, Fairness, and Opportunity in Algorithmic Decision-Making Systems. *Canadian Journal of Philosophy* 52, 1 (2022), 26–43. <https://doi.org/10.1017/can.2022.3>
- [13] Kathleen A. Creel. 2020. Transparency in Complex Computational Systems. *Philosophy of Science* 87, 4 (Oct. 2020), 568–589. <https://doi.org/10.1086/709729>
- [14] Jenny L Davis, Apryl Williams, and Michael W Yang. 2021. Algorithmic reparation. *Big Data & Society* 8, 2 (2021), 20539517211044808.
- [15] John Dewey. 1938. *Logic: The Theory of Inquiry*. (1938).
- [16] Benjamin Eidelson. 2021. Patterned Inequality, Compounding Injustice, and Algorithmic Prediction. *American Journal of Law and Equality* 1 (2021), 252–276.
- [17] Joseph Fishkin. 2014. *Bottlenecks: A new theory of equal opportunity*. Oxford University Press, USA.
- [18] Joseph Fishkin and William E Forbath. 2014. The anti-oligarchy constitution. *BUL Rev.* 94 (2014), 669.
- [19] Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2021. The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making. *Commun. ACM* 64, 4 (2021), 136–143.
- [20] Ben Green. 2022. Escaping the Impossibility of Fairness: From Formal to Substantive Algorithmic Fairness. *Philosophy and Technology* 35 (12 2022). Issue 4.
- [21] Ben Green and Salomé Viljoen. 2020. Algorithmic Realism: Expanding the Boundaries of Algorithmic Thought. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAccT* '20). Association for Computing Machinery, New York, NY, USA, 19–31. <https://doi.org/10.1145/3351095.3372840>
- [22] Mario Günther and Atoosa Kasirzadeh. 2021. Algorithmic and human decision making: for a double standard of transparency. *AI & SOCIETY* 37, 1 (April 2021), 375–381. <https://doi.org/10.1007/s00146-021-01200-5>
- [23] Moritz Hardt and Michael P Kim. 2022. Backward baselines: Is your model predicting the past? *arXiv preprint arXiv:2206.11673* (2022).
- [24] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016).
- [25] Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (Computationally-Identifiable) Masses. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1939–1948. <https://proceedings.mlr.press/v80/hebert-johnson18a.html>
- [26] Deborah Hellman. 2018. Indirect discrimination and the duty to avoid compounding injustice. *Foundations of Indirect Discrimination Law, Hart Publishing Company* (2018), 2017–53.
- [27] Anna Lauren Hoffmann. 2019. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society* 22, 7 (2019), 900–915.

- [28] Sune Holm. 2023. Algorithmic legitimacy in clinical decision-making. *Ethics and Information Technology* 25, 3 (July 2023). <https://doi.org/10.1007/s10676-023-09709-7>
- [29] Atoosa Kasirzadeh. 2022. Algorithmic Fairness and Structural Injustice: Insights from Feminist Political Philosophy. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 349–356.
- [30] Maximilian Kasy and Rediet Abebe. 2021. Fairness, equality, and power in algorithmic decision-making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 576–586.
- [31] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
- [32] Jon Kleinberg and Manish Raghavan. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118, 22 (2021), e2018340118. <https://doi.org/10.1073/pnas.2018340118> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2018340118>
- [33] Jon Kleinberg and Manish Raghavan. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118, 22 (2021), e2018340118.
- [34] Danielle Li, Lindsey R Raymond, and Peter Bergman. 2020. *Hiring as exploration*. Technical Report. National Bureau of Economic Research.
- [35] Martha Minow. 1990. *Making all the difference: Inclusion, exclusion, and American law*. Cornell University Press.
- [36] Brent Mittelstadt, Chris Russell, and Sandra Wachter. 2019. Explaining Explanations in AI. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM. <https://doi.org/10.1145/3287560.3287574>
- [37] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453. <https://doi.org/10.1126/science.aax2342> arXiv:<https://www.science.org/doi/pdf/10.1126/science.aax2342>
- [38] Cathy O’Neil. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- [39] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On Fairness and Calibration. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [40] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (*FAT* ’20*). Association for Computing Machinery, New York, NY, USA, 469–481. <https://doi.org/10.1145/3351095.3372828>
- [41] John Rawls. 2004. A theory of justice. In *Ethics*. Routledge, 229–234.
- [42] Reuters. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>.
- [43] MIT Technology Review. 2021. LinkedIn’s job-matching AI was biased. The company’s solution? More AI. <https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>.
- [44] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
- [45] Amartya Sen. 1980. *Equality of What?* Cambridge University Press, Cambridge.
- [46] Connor Toups, Rishi Bommasani, Kathleen A Creel, Sarah H Bana, Dan Jurafsky, and Percy Liang. 2023. Ecosystem-level Analysis of Deployed Machine Learning Reveals Homogeneous Outcomes. *arXiv preprint arXiv:2307.05862* (2023).
- [47] Suresh Venkatasubramanian and Mark Alfano. 2020. The philosophical basis of algorithmic recourse. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 284–293.
- [48] Ari Waldman and Kirsten Martin. 2022. Governing algorithmic decisions: The role of decision importance and governance on perceived legitimacy of algorithmic decisions. *Big Data & Society* 9, 1 (Jan. 2022), 205395172211004. <https://doi.org/10.1177/20539517221100449>
- [49] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2023. Against Predictive Optimization. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM. <https://doi.org/10.1145/3593013.3594030>
- [50] Bernard Williams. 1962. The Idea of Equality. In *Philosophy, Politics and Society*, Peter Laslett and Walter Runciman (Eds.). Blackwell, 112–117.
- [51] Indre Zliobaite. 2015. On the relation between accuracy and fairness in binary classification. *arXiv preprint arXiv:1505.05723* (2015).