

Fair Machine Learning Under Partial Compliance

Jessica Dai
Brown University
Providence, Rhode Island, USA
jessica.dai@alumni.brown.edu

Sina Fazelpour
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
sinaf@andrew.cmu.edu

Zachary C. Lipton
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
zlipton@cmu.edu

ABSTRACT

Typically, fair machine learning research focuses on a single decision maker and assumes that the underlying population is stationary. However, many of the critical domains motivating this work are characterized by competitive marketplaces with many decision makers. Realistically, we might expect only a subset of them to adopt any non-compulsory fairness-conscious policy, a situation that political philosophers call *partial compliance*. This possibility raises important questions: how does partial compliance and the consequent strategic behavior of decision subjects affect the allocation outcomes? If $k\%$ of employers were to voluntarily adopt a fairness-promoting intervention, should we expect $k\%$ progress (in aggregate) towards the benefits of universal adoption, or will the dynamics of partial compliance wash out the hoped-for benefits? How might adopting a global (versus local) perspective impact the conclusions of an auditor? In this paper, we propose a simple model of an employment market, leveraging simulation as a tool to explore the impact of both interaction effects and incentive effects on outcomes and auditing metrics. Our key findings are that at equilibrium: (1) partial compliance by $k\%$ of employers can result in far less than proportional ($k\%$) progress towards the full compliance outcomes; (2) the gap is more severe when fair employers match global (vs local) statistics; (3) choices of local vs global statistics can paint dramatically different pictures of the performance vis-a-vis fairness desiderata of compliant versus non-compliant employers; and (4) partial compliance based on local parity measures can induce extreme segregation. Finally, we discuss implications for auditors and insights concerning the design of regulatory frameworks.

CCS CONCEPTS

• **Social and professional topics** → **Governmental regulations; Socio-technical systems**; • **Applied computing** → **Law; Economics**; • **Computing methodologies** → **Modeling and simulation; Machine learning**.

KEYWORDS

fairness, distributive justice, hiring, simulations, segregation

ACM Reference Format:

Jessica Dai, Sina Fazelpour, and Zachary C. Lipton. 2021. Fair Machine Learning Under Partial Compliance. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21), May 19–21, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3461702.3462521>

1 INTRODUCTION

Responsible implementation of any allocation policy requires robust foresight about its likely impacts. In order to be useful, such an analysis needs to take into account existing and emerging interdependencies between the policy and environmental factors that shape the policy's long-term, situated consequences [22, 29]. However, to date, most studies of the performance and bias of algorithms applied to allocation decisions examine the algorithm in isolation, ignoring the wider deployment context. As a result, these analyses risk distorting our understanding of the impacts of specific algorithms, and limit our ability to anticipate broader societal implications of algorithmic decision-making.

Recently, a more critical thread in algorithmic fairness scholarship has called for a broader, systems-level approach to “fairness”, recognizing that algorithmic decisions do not happen in a vacuum [23, 27, 30, 32, 35, 41, 56]. Decisions may have long-term ramifications for individual welfare beyond the snapshot captured at the time of prediction [19, 41]. Thus, shifting attention towards the agency, impacts, and responsibility of decision makers in context is imperative.

In this paper, we adopt such a systems-level approach to explore the setting where multiple decision makers interact in a single labor market. Rather than considering the fairness of policies that a single decision maker might choose (i.e., the fairness of a single algorithm), we assume that there are several decision makers, whose decisions impact each another via market dynamics. While there are many possible settings in which a multi-decision maker scenario could take place—the provision of loans, for instance—we use the job market as a toy model for this scenario, both for simplicity and to set our work in dialogue with the broader labor economics literature addressing discrimination and partial compliance.

Two factors complicate the situation. First, employers vary in terms of their hiring policies, especially concerning their adherence to fairness-promoting measures. This situation of *partial compliance* reflects the current reality of predictive algorithms in hiring, which is characterized by heterogeneity across vendors regarding the type of measures, if any, enforced for counteracting bias [53]. Second, complicating matters further, differences in hiring policies across institutions can *incentivize* strategic applications, altering the distribution of candidates subsequently seen by employers [22].

We investigate these dynamics using simulation tools. Our models consist of two types of agents: applicants and employers. The

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIES '21, May 19–21, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8473-5/21/05...\$15.00

<https://doi.org/10.1145/3461702.3462521>

applicants each have a single “score” reflecting their perceived skill levels, and belong to one of two demographic groups: one which has been historically disadvantaged, and associated with lower scores on average, and one which has been historically advantaged, which has higher scores on average. In this work, we take no position on the extent to which this disparity is the result of systematic biases in the appraisal of the disadvantaged group, or is an accurate reflection of skills that vary because of upstream discrimination in society. Our general observations apply in both cases.

The employers may either be fairness-conscious (*compliant*)—taking into account considerations of demographic parity [13, 24], or fairness-agnostic (*non-compliant*)—deciding solely on the basis of scores. We also explore settings where applicants decide the type of institutions to which they apply strategically in light of the different incentives afforded by these selection policies.

We emphasize that our model is not intended as a realistic depiction of the labor market. We do not claim to offer direct policy prescriptions. Instead, our purpose is to propose the simplest conceivable model that captures the effects of partial compliance. By elucidating some basic qualitative insights regarding the impacts of partial compliance in allocative decisions, we aim to clarify the associated set of concerns that must be accounted for by any regulator. We argue that if even the most simple models evidence the complex interactive effects introduced by partial compliance behavior, then these effects must be considered when discussing the impact of specific policies or algorithmic approaches.

Our findings. Even with the simplest of assumptions, the relationships between the number of compliant institutions and various relevant metrics exhibit interesting phenomena:

- (1) Partial compliance (by $k\%$ of employers) can result in far less than proportional ($k\%$) progress towards the full compliance outcomes.
- (2) This gap (between the benefit of $k\%$ compliance and $k\%$ of the benefit of full compliance) is wider when compliant employers enforce demographic parity to match global (vs local) statistics.
- (3) Choices of global vs local statistics can paint dramatically different pictures of the performance of compliant (versus non-compliant) employers with respect to fairness considerations.
- (4) When coupled with incentive effects, partial compliance can induce extreme segregation across institutions.

Our results illuminate a critical shortcoming in current approaches to understanding fairness in algorithmic-based allocations, and have significant implications for how we think about auditing decision makers and assessing the potential benefits of regulation. For example, simulations with our model show that even if a large fraction of employers voluntarily comply with a fairness-promoting policy, that does not necessarily mean that a commensurate fraction of the benefit (relative to universal adoption) has been realized. Consequently, a regulator assessing the urgency of implementing fairness measure should take into account that even if only 20% of the population are non-compliant with a particular voluntary measure, they may be obstructing a much larger share, say 50% of the possible benefits of the policy. Moreover, our findings suggest that in order to understand an employer’s performance vis-a-vis

fairness desiderata, it is not enough to look at statistics calculated based on the stream of candidates that apply to them—we must also consider the way that the set of applicants that they encounter may diverge from the demographics of the general population, and how these dynamics involve both interactions among the employers and strategic behavior among applicants.

The rest of this paper is organized as follows: In Section 2, we survey literature from philosophy, (labor) economics, and the fair machine learning community, making connections to other work showing that the (partial) compliance among multiple decision makers is an essential consideration for assessing both moral responsibility and implementing practical measures. In Section 3, we introduce our model, including the parameters to our simulation and several axes of variation that we explore. In Section 4, we discuss our experiments and key results from those experiments. Finally, Section 5 provides a more critical discussion, including implications for regulating machine learning in allocative settings.

2 RELATED WORK

This work builds on several lines of research in economics, fair machine learning, political philosophy, and computational social science. An extensive literature in economics models discrimination in employment. Becker [10] introduced the notion of *taste-based discrimination*, where employers’ distaste for hiring employees from a certain group results in them behaving as though hiring a worker from the marginalized group was associated with a higher cost (a “disutility”), despite workers from both groups being identical in terms of true skill level. Becker also shows that this differential treatment among employers induces a sorting of minority employees into the least discriminatory employers, with the equilibrium wage determined by the disutility of the marginal discriminator. While our setup and motivation differ from Becker’s, with employers intervening to mitigate (rather than instigate) disparities, this segregation effect induced by differential treatment across employers also appears in our model.

Arrow et al. [5] famously criticized Becker’s model, arguing that discrimination thus characterized would decrease competitiveness and be driven out of the market, suggesting instead to focus on models of discrimination driven by imperfect information. Along these lines, Phelps [49] introduced a *statistical* model for discrimination in hiring, whereby disparities emerge due to differences in the difficulty of measuring the true skill level of each group of employees. Aigner and Cain [1] build on this idea, arguing that economic discrimination ought to be measured by differential treatment conditioning on true skill. By contrast, we take no position on whether observed scores accurately reflect the employee’s true skill level. Finally, Coate and Loury [18] address the long-term efficacy of affirmative-action policies, finding that, depending on specific parameter settings in their model, affirmative action can either eliminate stereotypes, or appear to confirm (untrue) negative stereotypes. As our “fairness intervened” models are functionally affirmative-action policies, we also explore the long-term dynamics of such policies. By contrast, we focus on the impacts of many employers adopting different policies on binary hiring decisions, not on concerns regarding stereotypes or wages.

Another related line of work calls for more realistic assumptions about the social context of allocation [23, 32, 41, 56]. In the fair machine learning literature, Hu and Chen [32] called attention to dynamics of employer-employee interactions, modeling the labor market as a series of principal-agent interactions. They draw upon the same threads of the economics literature, but focus on reputation and effort exertion. Liu et al. [41] focuses on credit ratings, showing that with a simple but reasonable set of assumed dynamics, certain fairness interventions can harm the very groups they are intended to protect. Hardt et al. [30], Hu et al. [33], Kilbertus et al. [37], Kleinberg and Raghavan [38], Milli et al. [44] all focus on the strategic behavior of individuals subject to automated decisions. Hu et al. [34] consider fairness in a setting where multiple classifiers interact with one another in the same system. Finally, Rambachan et al. [52] approaches fair machine learning from an economic perspective, constructing a social welfare function for a policymaker and a private objective function for an algorithm designer, investigating the relationship between disclosure and regulation. While these works recognize the problem of framing decisions as classifications, none focus on partial compliance, the central issue in this paper.

By contrast, we focus on two aspects of deployment dynamics that, though critical in shaping the ethical impact of algorithms in context, tend to be abstracted away in standard evaluations of algorithmic systems. First, our model represents potential differences among decision-makers in adherence to ethical or legal obligations, thus relaxing the assumption of a central decision-maker (or, equivalently, of full compliance), according to which all relevant agents comply with what justice demands of them. Present in many philosophical theories of justice and implicitly assumed by many works in fair machine learning [23], the full compliance assumption enables one to focus theorizing on the obligations that are the “fair share” of any agent. Nonetheless, recent philosophical works have cast doubt on whether theories developed under this assumption can provide sufficient practical guidance for agents in the actual world characterized by partial compliance [4, 62]. This line of work considers when and how in circumstances of partial compliance agents might face obligations that differ from what would have been their fair share, had others complied [43, 55, 62]. In the related labor economics literature, papers tend to focus on determining the incentive structures that promote or impede compliance with regulations such as minimum wage laws [6, 58], examining their macro-level impacts on the treatment of “non-favored” groups [16].

Second, in our models, decision subjects are represented as agents capable of responding strategically to the incentive structure of the environment. While abstracted away in most analyses of algorithmic reliability, this type of secondary effect is widespread in real-world allocation settings, and achieving foresight about its impacts is a priority for policy makers [22, 52]. Our work contributes to emerging efforts in the fair machine learning literature towards broadening the scope of analysis to include these effects [19, 30, 42]. Moreover, in exploring the impact of these dynamics, our work goes beyond assessments of algorithmic performance in static settings, furthering research on the long-term impact of proposed interventions [31, 32, 41].

We also build on recent research using simulation models to study fairness in ML systems [19]. While comparatively new in fair

machine learning, simulation studies represent a core methodology in economics and sociology [11, 15], and are increasingly used by philosophers to study social dynamics in general [66] and fairness in particular [45, 46]. Simulations are favored in these domains owing to their ability to model emergent outcomes of multiple interdependent decisions in non-stationary settings. Furthermore, particularly in the presence of heterogeneity in individual characteristics, simulations can yield insights that are not readily available in traditional aggregate models, such as those based on closed form solutions and/or systems of differential equations [36].

3 EXPERIMENTAL SETUP

We now provide a detailed description of the models explored in our simulations and motivate their design. In all of our models of a job market with partial compliance, all applicants have exactly two attributes: (i) a score, representing perceived skill for the job; and (ii) a group identity. Applicants may belong either to the advantaged group with higher mean score (Group A) or the marginalized group with lower mean score (Group B). Across our experiments, we consider two levels of representation in the broader population: one where the disadvantaged group constitutes 25% of the populations and another where they constitute 50%. Our market contains a number of employers ($n = 50$), $k\%$ of which may be compliant, and $(100 - k)\%$ of which are non-compliant.

At each time step, some number of new applicants ($a = 1250$) enter the job market. Each newcomer to the applicant pool is randomly assigned a group membership (according to population demographics). Each applicant’s score is drawn from a normal distribution: $\mathcal{N}(0, 1)$ for Group A, and $\mathcal{N}(0, -0.3)$ for Group B. Then, every applicant chooses one employer to apply to, and each institution hires $h = 10$ applicants. Once hired, applicants are removed from the market. Additionally, we remove applicants that have not been hired after 10 rounds.

3.1 How do institutions choose applicants?

We consider three possible policies that institutions may adopt when choosing applicants to hire: one generic non-compliant strategy, and two possible fairness-conscious (i.e. “compliant”) strategies, which satisfy some version of demographic parity.

- (1) *Generic policy.* Non-compliant employers simply hire the h highest-scoring applicants.
- (2) *Global parity policy.* Compliant employers with the global parity policy satisfy demographic parity with their hires, with respect to global demographics; this may or may not be the same as the demographics of their applicant pool. For example, if 25% of the population belonged to Group B, even if they accounted for 35% of applicants to a global-parity employer, they would only account for 25% of their hires.
- (3) *Local parity policy.* Compliant employers with the local parity policy satisfy demographic parity with respect to the demographics of their applicant pool at that round; in most cases, this is not the same as the overall demographics of the environment. For example, if 15% of applicants to a local-parity employer were from Group B, then 15% of the employer’s hires will be from Group B, even if Group B comprises 25% of the entire population.

The latter two parity strategies are probabilistic—hiring $x\%$ from Group B *in expectation*—rather than deterministically hiring a specific number from Group B based on a rounded proportion of available headcount. For simplicity, we only consider scenarios in which all compliant employers adopt the same strategy (either local or global).

Comments on demographic parity Our operationalization of fairness in terms of demographic parity is not intended as an endorsement of this measure as the appropriate fairness measure in hiring settings. Rather, our choice is based on the widespread use of the measure in current practice [53], perhaps due to a perceived connection between the quantitative measure and disparate impact doctrine in the United States [24] and indirect discrimination regulations in the European Union [2]¹.

Additionally, if available scores accurately reflect “true” skill level, then the generic, non-compliant policy may actually be fair according to some proposed definitions of fairness, such as calibration [50]. While our results are relevant regardless of the relationship between “true” and available scores, making this assumption means that our work can be also be re-interpreted as investigating the scenario where many intentionally-compliant employers have different *interpretations* of compliance—that is, employers are satisfying different definitions of fairness.

In the case that available scores *do not* accurately reflect “true” skill level for Group B, consider a setting where the true skill distributions are identical, and the compliant policy involves correction for the score difference rather than setting explicit headcounts. More concretely, the score-correcting compliant policy will simply add the known difference in scores to the scores of all applicants from Group B, then hire the h applicants with the highest (corrected) scores regardless of group membership—operationalizing fairness as treating individuals with the same true skill equally. As it turns out, this setting yields identical results to the local parity policy: for any given set of applications, a local parity employer will hire the top $x\%$ of applicants from each group. Meanwhile, a score-correcting employer corrects the scores of Group B, so that both groups have the exact same score distributions. Then, hiring the top $x\%$ based on corrected scores is equivalent to hiring the top $x\%$ from each group. However, we note that the two may diverge when applicants’ strategic behavior can be score aware.

Finally, we note that both possible compliant policies—local and global—are constrained by the demographics of the applicant pool, even in the global parity case: for example, 25% of Group B among all *applicants* may still reflect under-representation with respect to the entire population, which means that even a “global parity” employer only satisfies demographic parity with respect to the overall applicant pool, rather than the true global population demographics.

3.2 How do applicants choose institutions?

We also consider three possible strategies that applicants may employ when choosing institutions to apply to. Let $p_{G \in A, B}$ represent

the probability of an applicant from group G (either A or B) choosing to apply to a *compliant* institution, scaled by the total number of compliant institutions. Like the employer policies, these strategies are stochastic.

- (1) *Completely at random.* All applicants from both groups are equally likely to apply to institutions of either type; hence, $p_A = p_B = k$. This reflects *no* strategic behavior, i.e., applicants have no sensitivity to incentives.
- (2) *Static preference.* Over the course of the simulation, all applicants from Group A have a fixed preference for applying to a non-compliant institution, and all applicants from Group B have a fixed preference for applying to a compliant institution; hence, $p_A < k < p_B$. This reflects strategic behavior, where applicants have some knowledge about the nature of each institution’s policies, but no access to additional information over the course of the simulation—that is, applicants are sensitive to incentives but have limited knowledge of the system.
- (3) *Dynamic preference.* Over the course of the simulation, p_A and p_B are adjusted for each round based on the results of the previous round. For each group, if that group’s acceptance rate in *compliant* institutions is greater than its acceptance rate in *non-compliant* institutions, then the log odds ratio $\ln(p_G / (1 - p_G))$ is increased by a constant amount 0.05. Otherwise, it is decreased by the same amount. Equilibrium for each group is reached when the probability of being accepted at a parity institution is the same as the probability of being accepted at a generic institution. This reflects strategic behavior where applicants are aware of their group membership, have access to new information at each timestep, and are able to update their strategy accordingly.

Comments on applicant strategy and agent-based modeling

While we do not claim that these strategies exactly model the decision making processes of individuals in the real world, these coarse approximations of aggregate behavior yield valuable qualitative insights. Though we use a simple toy model, the core motivations for its design are grounded in reality. The hiring platform Applied, for example, claims that fairness-conscious hiring policies result in increased applications from minority groups [9]. It is impossible to exactly quantify the degree to which either applicant strategy (static or dynamic) represents “true” behavior. However, as mentioned in Section 2, simulation studies are a core methodology in both economics and philosophy, and in this work, the value of simulation is to test the qualitative impact of some sort of applicant strategy.

4 RESULTS

In all of our experiments, we vary the number of compliant institutions (out of 50 total) from 0 to 50. For each number of compliant institutions, we run ten trials of the simulation. For each trial, we run the simulation until it reaches equilibrium: 100 steps for static applicant strategy, and 200 steps for adaptive applicant strategy. We then continue running the simulation for the same number of additional timesteps and calculate statistics from each trial based on

¹See Lipton and Steinhardt [40] and Wachter et al. [63] for critical perspectives on the connection.

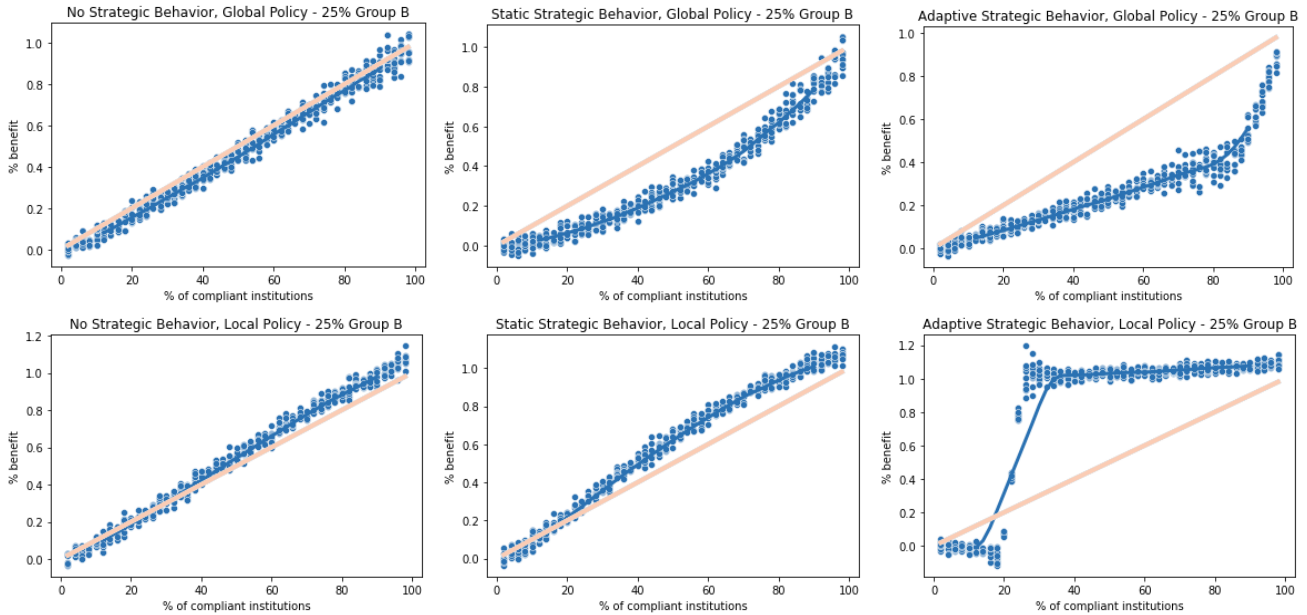


Figure 1: Benefit as measured by demographic parity for different institution policies. Far left plots show market where all applicants pick employers uniformly at random; as expected, we see exactly linear gain. In the center column, applicants have a slight preference for a more favorable employer (compliant for Group B, non-compliant for group A), and in the far right plots, applicants have an adaptive strategy.

the post-equilibrium timesteps. In all of our plots, one dot reflects the statistics calculated from a single trial.

Sublinear gain Our first key finding is that when employees apply strategically, then under partial compliance, the aggregate benefit from an additional compliant employer depends strongly on how many institutions are already compliant. In Figure 1, all employees apply with the strategy of static preference: that is, knowing that compliant employers are more likely to hire Group B applicants, and that non-compliant employers are more likely to hire Group A applicants, employees from Group B apply to compliant employers with probability 0.55 (scaled by number of each type of employer) and employees from group A apply to non-compliant employers with probability 0.55. The y-axis is scaled demographic parity, where $y = 0$ corresponds to the disparate impact score $\frac{P(\text{hired}|B)}{P(\text{hired}|A)}$ when all employers are non-compliant (with our main experimental parameters, this is approximately 0.75), and $y = 1$ corresponds to “perfect” parity. One might hope that $k\%$ compliance would correspond to at least $k\%$ of the benefits, a condition that we denote *linear gain*. In Figures 1 and 2, this is illustrated by the light peach line.

Notably, when all compliant institutions satisfy fairness with respect to *global* statistics, the partial compliance curve is convex, illustrating *sublinear gain*— $k\%$ compliance always gives less than $k\%$ of the attainable benefit. Perhaps this should not be a surprising result. The baseline demographic parity (% benefit = 0, at 0% compliance) reflects a scenario where each (non-compliant) employer receives an applicant pool that reflects the overall demographics

of the system (i.e., if 25% of all applicants in the system are from Group B, then on average 25% of non-compliant employers’ applicants also are from Group B). In order for *linear gain* to occur, then at $k\%$ compliance, all $(100 - k)\%$ non-compliant employers must hire at the same rate as they were at 0% compliance even as the $k\%$ compliant employers hire exactly in accordance with global demographic parity. However, due to applicant strategy, the distribution of applicants to non-compliant employers at $k\%$ compliance no longer reflects the demographics of the system. Instead, non-compliant employers see relatively more Group A applicants and relatively fewer Group B applicants. As a result, the non-compliant hiring strategy results in an even lower percentage of Group B (as a proportion of overall hires) than at 0% compliance, giving rise to sublinear gain.

Under local parity policies, the partial compliance curve can actually reflect *superlinear* gain, as when Group B constitutes 25% of the population. However, when Group B constitutes 50% of the population (Figure 2), these dynamics change: local parity policies now also induce *sublinear* gain, and the global parity curve indicates a more pronounced sublinear gain.

Regardless of whether Group B comprises 25% or 50% of the population, following the *global* parity policy leads to comparatively worse gains than following the *local* parity policy—that is, for any given $k\%$ compliant institutions, the percent benefit when employers satisfy global parity is lower than when employers satisfy local parity. The explanation, both for super/sub-linearity of local parity policies, and for why sublinear gain under global parity is always worse than under local parity, lies in the flexibility that a local parity policy affords. Under global parity policies, the fraction of hires

that a compliant institution can make from Group B is fixed based on their share of the underlying population. However, with local parity policy, it is possible for all $k\%$ of the compliant employers to allocate their entire headcount to Group B (in the event that Group B comes to represent 100% of their applicants). Thus, under local parity, compliant employers are able to take on more than their “fair share” (to borrow terminology from the philosophy literature on partial compliance).




Figure 2: Aggregate statistics when Group B is 50% of the population; benefit is measured by overall demographic parity. Left: static applicant strategy; right: adaptive applicant strategy.

Static vs adaptive applicant strategy When employees are able to update their application strategy at each timestep, interesting dynamics emerge (Figure 1, 2). Recall that the likelihood of employees from a given group applying to each type of employer (compliant vs non-compliant) is adjusted based on group-wise acceptance rates from the previous timestep. Hence, equilibrium for each group is reached when that group encounters the same acceptance rate from both compliant and non-compliant employers. Under global parity policies, the first 80% of compliant institutions are only able to push the macro-level statistics around *halfway* to parity; the remaining 50% of benefits relies entirely on the last 20% of employers becoming compliant. Interestingly, under local parity policies when Group B is 25% of the population, the first 20% of compliant employers have functionally no effect on the macro-level view of fairness, while complete parity is achieved by the time around 30% of employers are compliant. For intuition as to why this is the case, we can look at the equilibrium probabilities for applying to either type of employer. Figure 3 shows that under local parity policies, the equilibrium p_B (probability of applying to a compliant employer for Group B) quickly goes to 1. With 20% or more compliant institutions, Group B always applies almost exclusively to compliant institutions. Meanwhile, until 26% or more employers are compliant, Group A applies almost exclusively to non-compliant




Figure 3: Groupwise equilibrium probability (p_A and p_B described in Section 3) of applying to either compliant or non-compliant employers, under adaptive applicant strategy. The orange line indicates the p reflecting no preference (i.e. probability determined solely by the proportion of compliant institutions currently in the system, $p = k$). Left: global parity policy; right: local parity policy.

institutions. Under global parity policies, the difference in preference induced by partial adoption of the fairness-promoting policy is less severe.

The emergent demographic composition of institutions A closer look at institution-specific outcomes reveals that at equilibrium, strategic applications can result in *homogeneity* within institutions and *segregation* across institutions. In the case of global parity policies, the dramatic increase in aggregate parity (Figure 1, right column) is coupled with a precipitous drop-off in the percentage of hired applicants belonging to Group B in non-compliant institutions (Figure 4, bottom left). The situation is even more dire under local parity policies, as the the equilibrium strategies mean that non-compliant institutions have *no* hired applicants (or indeed, applications) from members of Group B (Figure 4, bottom right). Notably, though the aggregate parity curves under the global policy do not look so different in Figure 1, the segregation effects do *not* occur when applicants operate under a *static* application strategy: while partial compliance has some impact on the overall demographic composition of hired employees, the percentage of Group B never approaches zero (Figure 4, top row).

The impact of the original demographic makeup on adaptive applicant strategy When employees were applying to firms under a *static* strategy, the impact of changing from a scenario where Group B is 25% of the population (Figure 1) to one where Group B is 50% of the population (Figure 2), while significant, affects aggregate statistics in similar ways at all levels of compliance and for both global and local parity policies. However, when applicant strategies are adaptive, increasing the proportion of Group B in the population (Figure 2) means that under global parity policies, the first 80% of compliant institutions—despite reaching 50% of the benefit when Group B was 25% of the population (Figure 1)—actually have *no* impact on aggregate demographic parity. The critical tipping point, however, remains the same, at 80% compliance. Under local parity policies, on the other hand, the overall shape of the aggregate parity curve remains the same—two large regions with either zero or perfect parity, and one small intermediary

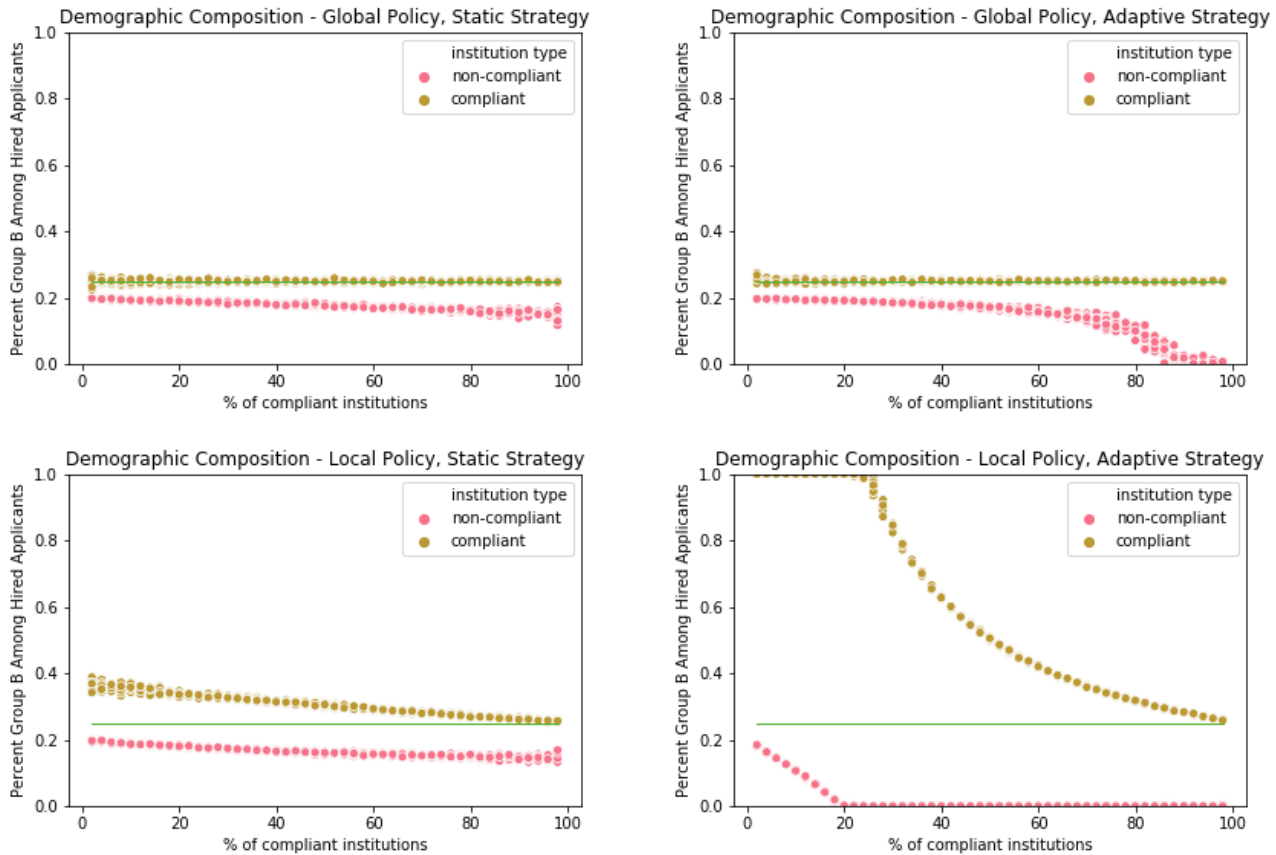


Figure 4: Demographic composition among hired employees, by institution type—in these graphs, Group B is 25% of the population. Top row: Applicants employ a *static* strategy. Bottom row: applicants employ an *adaptive* strategy. Light green horizontal line indicates percentage of Group B in population.

transition region—but when Group B comprises 50% of the population, the critical transition region is between 40%-50% compliance, rather than 20%-30% compliance.

5 DISCUSSION

Our simulations illustrate several fundamental but commonly overlooked issues that plague the ethical evaluation and governance of algorithmic tools in consequential allocation settings. While our results do not imply specific or prescriptive policy solutions, they do raise critical questions about the design and adoption of fair policies.

Beyond narrow assessments of fairness: diversity and integration Consider first that, in many allocative contexts, task-related utility and fairness are not the only desiderata. For example, in hiring contexts, diversity within the workforce is intrinsically valuable, both due to its potential to enhance team performance and on moral and political grounds [47, 59]. While recent work in fair ML has begun to consider the interaction between diversity, utility and fairness [14, 21], most analyses remain restricted to

static settings, focused on individual decision-makers, neglecting the interactions among their decisions and those of their peers and the influence of dynamic factors, such as incentive effects, on long-term policy consequences. Consider what Steel et al. [60] refer to as the representative concept of diversity (see also Smith-Doerr et al. [57]), motivated by concerns about democratic legitimacy, which requires the distributional properties of the selected group to match those of the general population. The global demographic parity measure thus tracks this notion of diversity. Viewed through a static lens, and setting aside the influence of incentives on the choice behavior of applicants, the same connection could be said to hold between the diversity concept and local demographic parity measures. Indeed, this has led some authors to roughly equate these notions of diversity and fairness [14]. The situation becomes more complicated, however, once the dynamics of adaptive application are taken into account. Here, the appearance of (ostensibly desirable) parity at the aggregate level conceals the detrimental impact of local parity policies on diversity within the workforces of the individual employers. These outcomes can emerge absent any explicit desire for segregation on the part of applicants or employers; rather, they are a consequence of the dynamics of incentive effects




Figure 5: Percent hired per group per institution type. Group B is 25% of the overall population.

under partial compliance. In addition to stripping institutions of the benefits of diversity, the resulting segregation can exacerbate the homophily-based processes that, according to a number of authors [3, 46], can cultivate or amplify injustice.

The aims and the value-alignment of regulation The above discussion indicates the urgent need to clarify the aims and value orientation of regulation. As Rambachan et al. [52] note, many discussions of regulation related to algorithmic fairness are fundamentally concerned with selecting a policy that will generate an optimal distribution of outcomes. Naturally, this requires first deciding what constitutes the optimal outcome distribution.

It is useful to frame this issue by inquiring about the aims of the policy that might support the enforcement of local (vs global) demographic parity. In practice, demographic parity is popular, perhaps owing to the 80% rule, which is sometimes invoked as a statistical test in the first phase of disparate impact cases [24]. Note, however, that this connection does not provide a blind endorsement of this form of parity as that which ought to be enforced. Certainly, demographic parity can be a part of a *diagnostic toolbox*, serving to indicate disparities that *could*, but need not, indicate underlying discrimination [8, 39]. When precisely measured, demographic disparity can signal moral or legal failings with that particular employer which lie outside the narrow scope of the quantitative measure itself. However, even when the disparity *is* a symptom of underlying ethical troubles with an allocation policy, enforcing the measure may be a misguided remedy to addressing these troubles (e.g., when the trouble lies with the choice of target outcomes or labels).

Another way of motivating the enforcement of (some form of) demographic parity is by reference to an employer’s wish to implement *affirmative action*. That is, employers may wish to enforce demographic parity, and so preferentially select applicants on the basis of their group membership, as a means of complying with a

moral obligation to increase the representation of historically disadvantaged social groups in their institutions. This interpretation resonates with the suggestions that, in some cases, the use of measures such as demographic parity is motivated by the “long-term societal goal” of living in a society where protected attributes are independent of task-relevant outcomes [7]. However, specifying the relation between demographic parity and affirmative action requires clarity about the underlying aim and justifications of the latter—issues that vary radically across different models of affirmative actions [3]—and considerations of whether the former indeed serves those aims. Crucially, our results indicate that, even with minimal incorporation of deployment dynamics, the (partial) adoption of local demographic parity is inconsistent with prominent *future-oriented* justifications of affirmative action. In particular, the emergence of between-institute segregation and a lack of within-institute diversity in our simulations indicate that partial compliance with the measure can result in significant conflicts with diversity-based [25] and integration-based [3] arguments for affirmative action.

Of course, one could adopt a different model of affirmative action to motivate the enforcement of demographic parity. For instance, depending on the interpretation of scores in our model (e.g., as a result of past, upstream injustices, or as an outcome of ongoing biases in an employer’s hiring practices), the measure could be connected to compensation-based [61] or discrimination-offsetting [64] justifications. Each of these models faces its own set of challenges, including discordance with the actual practice of law, failure to account for the *weight* given to social categories in preferential selection, engendering the expressive harm of *stigmatization*, and undermining the societal legitimacy of affirmative action [3, 25].

While adjudicating between different models of affirmative action is beyond the scope of this paper, it raises an important concern: Decisions about the aims and the alignment of regulation are value-laden to their core. As a result, these decisions should be made transparently, and on the basis of an integrated consideration of the relevant moral and political models. Importantly, our results show that individual efforts (or the lack thereof) to promote fairness can remain out of sight unless assessed through a more comprehensive, dynamic lens. Analyses of the kind carried out in this paper can not only bring these value judgments into the open, but also complement theorizing about which moral and political models we should prefer, and why. For example, analyses of deployment dynamics can offer qualitative insights about other meaningful endpoints and value-relevant considerations (e.g., diversity) that are likely to be relevant to assessing the desirability of alternative policies in context. Such approaches can thus contribute to recent calls for enriching the evaluation of downstream impacts of algorithmic decision-making [48]. Viewed from this perspective, far from simply being a burden to be neglected in the context of ethical design, assessment of deployment dynamics can guide our deliberations in such contexts.

Partial compliance and the design of appropriate auditing frameworks The type of partial compliance explored in this paper is a simple representation of the kinds of heterogeneity that exist in the adoption of fairness-promoting measures among various employers both in the use of algorithmic tools in hiring [53]

and in hiring more generally. The varied choices of measures is a *consequence* of the ambiguity of current regulatory frameworks. Indeed, the laxity of constraints provides even the *non-compliant* employers in our simulations with a claim to fairness. That is, insofar as these employers have access to the “ground truths” for skill scores, they can be seen as employing a perfect predictor that satisfies a number of other fairness desiderata suggested in the literature, such as parities in sensitivity, specificity, and precision across groups [17, 20]. Similar to the evaluative practices that inform them, these regulatory frameworks appear to be based on unrealistic assumptions and abstractions of the problem domain.

Our exploration of the dynamics of partial compliance raises central concerns that should inform judgments about both the need for regulation and the form that it should take. The discussion above relies heavily on the assumption that a regulator would be able to bring about something approximating full compliance to begin with. Determining how the behavior of individual decision makers compares to the behavior of all decision makers—and by extension, whether partial compliance is occurring—is therefore a critical concern for any regulatory regime.

Existing approaches to auditing have focused on examining the performance of a single algorithmic decision maker [26, 51, 54]. Similarly, Wilson et al. [65]’s work with Pymetrics, a hiring platform that uses local demographic parity, explicitly considers only the pool of applicants that Pymetrics receives, seeking to verify the extent to which the selection procedure adheres to this version of demographic parity.²

However, in addition to highlighting the potential cost to diversity and integration, our analysis shows that fairness statistics reported by each employer are impacted not only by their policies, but also by those of their competitors. In our simulation, at equilibrium under partial compliance, when employers adopt the global parity policy, an auditor looking at the fractions of applicants from Group A and Group B hired might erroneously conclude that compliant and non-compliant employers were behaving similarly (Figure 5). However, this mistaken view fails to account for the incentive effects, whereby compliant employers come to receive many more applications from members of the disadvantaged group. Thus, when auditing performance vis-a-vis ethical desiderata, we may not be able to determine how a firm is performing without also evaluating their peers.

Phenomena of this sort are not exclusive to partial compliance settings. D’Amour et al. [19], which study the long-term impact of (fair) decisions, find a similar instance of Simpson’s paradox where enforcing equal opportunity at each point in time does not result in equal opportunity in the aggregate, in the presence of interactive effects between decisions and the characteristics of decision subjects. Taken together, these results suggest a need for auditors to investigate not only the distributions of outcomes given the data, but the actual underlying policy. To this end, Rambachan et al. [52], who study the construction of ideal (fair) policy from the perspective of a regulator, find a particularly interesting result: the ideal disclosure regime is one where individual decision makers must disclose

²Our emphasis here is on the fact that the audit is solely focused on Pymetrics, not to claim that Pymetrics functions as a decision maker in the same way as employers do in our simulation. Additional discussion of the Pymetrics audit’s coverage, while merited, is out of scope for this work.

all information about their algorithm and decision rule, and the effectiveness of regulation is substantially diluted when disclosure is more limited. Finally, although the downstream consequences of *regulation* in a dynamic environment is beyond the scope of this work, viewing regulation under a dynamic lens suggests that the potential for partial compliance to mask the efforts of compliant institutions may provide an incentive for those institutions to share information about their policies with auditors despite the desire to protect proprietary information, because it may help differentiate themselves from non-compliant institutions.

In some sense, the abstractions in our model *underestimate* the implications of partial compliance for current regulatory and evaluative practices. This is because our model represents partial compliance only with respect to *concurrent* policies in a competitive marketplace of hiring. That is, we do not consider allocations that are *upstream* (e.g., in education) and *downstream* (e.g., promotion, mobility across work sector, banking) from hiring decisions, each made by decision-makers who may or may not adhere to their legal (or moral) obligations.

Elster [22] makes vivid the significance of such allocations for the well-being and opportunities of individuals:

The life chances of the citizen in modern societies ... depend on allocations made by relatively autonomous institutions, beginning with admission or nonadmission to nursery school and ending with admission or nonadmission to nursing homes. One could write the fictional biography of a typical citizen, to depict his life as shaped by successive encounters with institutions that have the power to accord or deny him the scarce goods that he seeks [22, p. 2].

Despite the potential of unexpected outcome due to robust couplings between policies at successive allocative settings, the implications of partial compliance at successive stages remain under-investigated by the fair ML community. This is a challenge that requires a concerted interdisciplinary effort by the community. Indeed, providing practical guidance under partial compliance poses a challenge to traditional frameworks of distributive justice in political philosophy. While looking to these frameworks for robust conceptual underpinnings of fairness measures can be fruitful [12], they were mainly concerned with modeling the re-distributive obligations of a nation state towards its citizens from the perspective of economic justice. However, when our focus is to provide guidance to relatively autonomous decision-makers using ML tools in local allocative settings, we can no longer simply operate with the same assumptions. Responsible innovation in general [28] and ethical deployment of algorithmic-based decision-making in particular [23] require more comprehensive foresight studies that are equipped to deal with the complexities of the deployment context. We hope that our work contributes a few preliminary steps towards this aim.

ACKNOWLEDGMENTS

Zachary Lipton thanks the Block Center for Technology and Society, Amazon AI, and NSF: Fair AI Award IIS2040929 for supporting ACMI lab’s research on the responsible use of machine learning. ZL is also grateful to PwC USA for funding this research through the Digital Transformation and Innovation Center sponsored by

PwC. Sina Fazelpour thanks the Block Center for Technology and Society and the Social Sciences and Humanities Research Council of Canada (award number 756-2019-0289). Jessica Dai was supported in part by a LINK award from Brown University.

REFERENCES

- [1] Dennis J Aigner and Glen G Cain. 1977. Statistical theories of discrimination in labor markets. *Ilr Review* 30, 2 (1977), 175–187.
- [2] Andrew Altman. 2020. Discrimination. In *The Stanford Encyclopedia of Philosophy* (summer 2020 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [3] Elizabeth Anderson. 2010. *The Imperative of Integration*. Princeton University Press, Princeton.
- [4] Kwame Anthony Appiah. 2017. *As If: Idealization and Ideals*. Harvard University Press, Cambridge.
- [5] Kenneth Arrow et al. 1973. The theory of discrimination. *Discrimination in labor markets* 3, 10 (1973), 3–33.
- [6] Orley Ashenfelter and Robert S Smith. 1979. Compliance with the minimum wage law. *Journal of Political Economy* 87, 2 (1979), 333–350.
- [7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [8] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [9] Be Applied Ltd. 2020. Why Applied? <https://www.beapplied.com/why-applied>
- [10] Gary S Becker. 1957. The economics of discrimination Chicago. *University of Chicago* (1957).
- [11] Federico Bianchi and Flaminio Squazzoni. 2015. Agent-based models in sociology. *Wiley Interdisciplinary Reviews: Computational Statistics* 7, 4 (2015), 284–306.
- [12] Reuben Binns. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on Fairness, Accountability and Transparency*. 149–159.
- [13] Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. 2009. Building Classifiers with Interdependency Constraints. In *Proceedings of the 2009 IEEE International Conference on Data Mining Workshops (ICDMW '09)*. IEEE Computer Society, Washington, DC, USA, 13–18.
- [14] L Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K Vishnoi. 2016. How to be fair and diverse? *arXiv preprint arXiv:1610.07183* (2016).
- [15] Damon Centola. 2018. *How behavior spreads: The science of complex contagions*. Vol. 3. Princeton University Press.
- [16] Yang-Ming Chang, Bhavneet Walia, et al. 2007. Wage discrimination and partial compliance with the minimum wage law. *Economics Bulletin* 10, 4 (2007), 1–7.
- [17] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. In *Big Data*. Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA.
- [18] Stephen Coate and Glenn C Loury. 1993. Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review* (1993), 1220–1240.
- [19] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 525–534.
- [20] William Dieterich, Christina Mendoza, and Tim Brennan. 2016. COMPAS risk scales: Demonstrating accuracy equity and predictive parity. *Northpointe Inc* (2016).
- [21] Marina Drosou, HV Jagadish, Evaggelia Pitoura, and Julia Stoyanovich. 2017. Diversity in big data: A review. *Big data* 5, 2 (2017), 73–84.
- [22] Jon Elster. 1992. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation.
- [23] Sina Fazelpour and Zachary C Lipton. 2020. Algorithmic Fairness from a Non-ideal Perspective. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 57–63.
- [24] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 259–268.
- [25] Robert Fullinwider. 2018. Affirmative Action. In *The Stanford Encyclopedia of Philosophy* (summer 2018 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [26] Ben Green and Yiling Chen. 2019. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 90–99.
- [27] Ben Green and Lily Hu. 2018. The myth in the methodology: Towards a recontextualization of fairness in machine learning. In *Proceedings of the machine learning: the debates workshop*.
- [28] Armin Grunwald. 2014. Technology assessment for responsible innovation. In *Responsible Innovation 1*. Springer, 15–31.
- [29] S Hansson. 2013. *The ethics of risk: Ethical analysis in an uncertain world*. Springer.
- [30] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. 2016. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*. 111–122.
- [31] Hoda Heidari, Vedant Nanda, and Krishna P Gummadi. 2019. On the long-term impact of algorithmic decision policies: Effort unfairness and feature segregation through social learning. *arXiv preprint arXiv:1903.01209* (2019).
- [32] Lily Hu and Yiling Chen. 2018. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*. 1389–1398.
- [33] Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. 2018. The Disparate Effects of Strategic Manipulation. In *Conference on Fairness Accountability and Transparency (FAT*)*.
- [34] Yaowei Hu, Yongkai Wu, Lu Zhang, and Xintao Wu. 2020. Fair Multiple Decision Making Through Soft Interventions. *Advances in Neural Information Processing Systems* 33 (2020).
- [35] Maximilian Kasy and Rediet Abebe. 2020. *Fairness, equality, and power in algorithmic decision making*. Technical Report. Working paper.
- [36] Elmar Kiesling, Markus Günther, Christian Stummer, and Lea M Wakolbinger. 2012. Agent-based simulation of innovation diffusion: a review. *Central European Journal of Operations Research* 20, 2 (2012), 183–230.
- [37] Niki Kilbertus, Manuel Gomez Rodriguez, Bernhard Schölkopf, Krikamol Muandet, and Isabel Valera. 2020. Fair decisions despite imperfect predictions. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 277–287.
- [38] Jon Kleinberg and Manish Raghavan. 2019. How Do Classifiers Induce Agents to Invest Effort Strategically?. In *ACM Conference on Economics and Computation (EC)*.
- [39] Zachary Lipton, Julian McAuley, and Alexandra Chouldechova. 2018. Does mitigating ML's impact disparity require treatment disparity?. In *Advances in Neural Information Processing Systems*. 8125–8135.
- [40] Zachary C. Lipton and Jacob Steinhardt. 2018. Troubling Trends in Machine Learning Scholarship. *Communications of the ACM (CACM)* 62, 6 (2018), 45–53.
- [41] Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2019. Delayed impact of fair machine learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. AAAI Press, 6196–6200.
- [42] Lydia T Liu, Ashia Wilson, Nika Haghtalab, Adam Tauman Kalai, Christian Borgs, and Jennifer Chayes. 2020. The disparate equilibria of algorithmic decision making when individuals invest rationally. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 381–391.
- [43] David Miller. 2011. Taking Up the Slack? Responsibility and Justice in Situations of Partial Compliance. In *Responsibility and Distributive Justice*, Carl Knight and Zofia Stemplowska (Eds.). Oxford University Press, 230–45.
- [44] Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. 2018. The Social Cost of Strategic Classification. In *Conference on Fairness Accountability and Transparency (FAT*)*.
- [45] Ryan Muldoon. 2016. *Social contract theory for a diverse world: Beyond tolerance*. Taylor & Francis.
- [46] Cailin O'Connor. 2019. *The Origins of Unfairness*. Oxford University Press, Oxford.
- [47] Scott E Page. 2019. *The diversity bonus: How great teams pay off in the knowledge economy*. Princeton University Press.
- [48] Ravi B Parikh, Ziad Obermeyer, and Amol S Navathe. 2019. Regulation of predictive analytics in medicine. *Science* 363, 6429 (2019), 810–812.
- [49] Edmund S Phelps. 1972. The statistical theory of racism and sexism. *The american economic review* 62, 4 (1972), 659–661.
- [50] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On Fairness and Calibration. *Advances in Neural Information Processing Systems* 30 (2017), 5680–5689.
- [51] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 33–44.
- [52] Ashesh Rambachan, Jon Kleinberg, Sendhil Mullainathan, and Jens Ludwig. 2020. *An economic approach to regulating algorithms*. Technical Report. National Bureau of Economic Research.
- [53] Javier Sánchez-Monedero, Lina Dencik, and Lilian Edwards. 2020. What does it mean to 'solve' the problem of discrimination in hiring? social, technical and legal perspectives from the UK on automated hiring systems. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 458–468.
- [54] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry* 22 (2014), 4349–4357.
- [55] Tamar Schapiro. 2003. Compliance, complicity, and the nature of nonideal conditions. *The Journal of Philosophy* (2003).
- [56] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Fairness, Accountability, and Transparency (FAT*)*.

- [57] Laurel Smith-Doerr, Sharla N Alegria, and Timothy Sacco. 2017. How diversity matters in the US science and engineering workforce: A critical review considering integration in teams, fields, and organizational contexts. *Engaging Science, Technology, and Society* 3 (2017), 139–153.
- [58] Lyn Squire and Sethaput Suthiwart-Narueput. 1997. The Impact of Labor Market Regulations. *The World Bank Economic Review* (1997), 119–143.
- [59] Daniel Steel, Sina Fazelpour, Bianca Crewe, and Kinley Gillette. 2019. Information elaboration and epistemic effects of diversity. *Synthese* (2019), 1–21.
- [60] Daniel Steel, Sina Fazelpour, Kinley Gillette, Bianca Crewe, and Michael Burgess. 2018. Multiple diversity concepts and their ethical-epistemic implications. *European Journal for Philosophy of Science* 8, 3 (2018), 761–780.
- [61] Judith Jarvis Thomson. 1973. Preferential hiring. *Philosophy & Public Affairs* (1973), 364–384.
- [62] Laura Valentini. 2012. Ideal vs. Non-ideal Theory: A Conceptual Map. *Philosophy Compass* (2012).
- [63] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2020. Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. *Available at SSRN* (2020).
- [64] Mary Anne Warren. 1977. Secondary sexism and quota hiring. *Philosophy & Public Affairs* (1977), 240–261.
- [65] Christo Wilson, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. [n.d.]. Building and Auditing Fair Algorithms: A Case Study in Candidate Screening. ([n. d.]).
- [66] Kevin JS Zollman. 2013. Network epistemology: Communication in epistemic communities. *Philosophy Compass* 8, 1 (2013), 15–27.