
On the Long-term Impact of Algorithmic Decision Policies: Effort Unfairness and Feature Segregation through Social Learning

Hoda Heidari^{*1} Vedant Nanda^{*2} Krishna P. Gummadi²

Abstract

Most existing notions of algorithmic fairness are *one-shot*: they ensure some form of allocative equality at the time of decision making, but do not account for the *adverse impact* of the algorithmic decisions today on the *long-term* welfare and prosperity of certain segments of the population. We take a broader perspective on algorithmic fairness. We propose an *effort*-based measure of fairness and present a data-driven framework for characterizing the long-term impact of algorithmic policies on reshaping the underlying population. Motivated by the psychological literature on *social learning* and the economic literature on *equality of opportunity*, we propose a micro-scale model of how individuals may respond to decision making algorithms. We employ existing measures of segregation from sociology and economics to quantify the resulting macro-scale population-level change. Importantly, we observe that different models may shift the group-conditional distribution of qualifications in different directions. Our findings raise a number of important questions regarding the formalization of fairness for decision-making models.

1. Introduction

Machine Learning tools are increasingly employed to make consequential decisions for human subjects, in areas such as credit lending (Petrasic et al., 2017), policing (Rudin, 2013), criminal justice (Barry-Jester et al., 2015), and medicine (Deo, 2015). Decisions made by these algorithms can have a long-lasting impact on people’s lives and may affect certain individuals or social groups negatively (Sweeney, 2013; Angwin et al., 2016; Levin, 2016).

^{*}Equal contribution ¹Computer Science Department, ETH Zürich, Zürich, Switzerland ²MPI-SWS, Saarbrücken, Germany. Correspondence to: Hoda Heidari <hheidari@inf.ethz.ch>.

This realization has recently spawned an active area of research into quantifying and guaranteeing fairness for machine learning (Dwork et al., 2012; Kleinberg et al., 2017; Hardt et al., 2016).

Most existing notions of fairness assume a *static* population: they ensure some form of *allocative equality* at the time of decision making, but do not account for the *adverse impact* of algorithmic decisions today on the *long term* welfare and prosperity of different segments of the population. For instance, consider equality of odds (Hardt et al., 2016). The notion requires that the model distributes different types of error (i.e., false positives and false negatives) equally across different social groups. But it does not take into consideration the fact that for members of the advantaged group these erroneous predictions may be easy to overturn, whereas for the disadvantaged it may take a significant amount of *effort* to improve their qualifications to obtain better algorithmic outcomes. Furthermore, in the long run, the decision-making model may nudge different segments of the population to obtain very different sets of qualifications—some of which might be socially and economically more desirable than others. This may in effect lead to further *marginalization* of these groups.

Motivated by these concerns about existing notions of fairness, we argue for a broader view of algorithmic models—one that treats them as *policies* implemented within a social context and with the potential of impacting individuals and reshaping society. Among other considerations¹, such view of decision-making models necessitates a deeper understanding of how individual decision subjects may *respond* to these models and how those responses may translate into *adverse impact* for certain segments of the population.

In this work, we propose an effort-based measure of unfairness for algorithmic decisions. We define a data-driven, group-dependent measure of *effort* drawing on the economic literature on Equality of Opportunity (Roemer & Trannoy, 2015). Our effort function captures the idea that the kind of

¹Another important factor is how a utility maximizing decision maker—employing the model—would respond to its predictions. For instance, they may interpret the predictions in a certain way, or update the model entirely. Prior work (Liu et al., 2018; Kannan et al., 2019) has already addressed some of these considerations.

changes required to obtain a desirable algorithmic outcome (e.g., changing one’s school type from public to private to get a better prediction for SAT score) is often significantly more difficult to make for members of the disadvantaged group compared to the advantaged. Building on this notion of effort, we formulate *effort unfairness* as the inequality in the amount of effort required for members of each group to obtain their desired outcomes.

To formulate the long-term impact of algorithmic policies on the underlying population, we specify a micro-scale model of how individuals respond to algorithmic decision-making models, taking inspiration from the psychological literature on *social learning* (Bandura, 1962; 1978). We posit that individuals observe and imitate the qualifications of their *social models*—someone who has received a better algorithmic outcome from the decision-making model—if by doing so, they can obtain higher rewards (Bandura, 1962; Apesteguia et al., 2007). More precisely, we model an individual’s response to be the decision-making algorithm by first selecting a social model for him/her; the individual is then assumed to exert effort to attain his/her model’s qualifications if and only if doing so improves his/her overall utility. With this individual-level behavioral model in place, we can simulate decision subjects’ responses and quantify the macro-scale impact of algorithmic policies on reshaping the underlying populations. We employ existing measures of segregation from sociology and economics (Massey & Denton, 1988) to characterize how the distribution of qualification for each group changes in response to the deployed model. Importantly, we observe that different models may shift the group-conditional distribution of qualifications in vastly different directions.

Our work raises a number of important questions about algorithmic policies and the formulation of fairness: What is the ultimate purpose of a fair predictive model—to guarantee allocative equality today, or to ensure similar distributions of qualifications in the long-run? With respect to short term allocative equality, are all error created equal, or should we take into account the disparity in the effort it takes for different groups to obtain their desired predictions? In the long run, what are the type of changes that different predictive models impose on the society? Which ones are desirable, and which ones should be watched out for? Is it ethically and economically acceptable to nudge different segments of the population toward obtaining different qualifications? If not, how can we prevent this without employing a model whose decisions may be perceived as unfair today? These are all critical questions that must be carefully analyzed before determining which model is fair and best-suited to make consequential decisions for humans. Addressing such ethical challenges is outside the scope of this paper—and arguably intradisciplinary Machine Learning research. We hope that our work serves as a reminder to the ML com-

munity that to formalize fairness appropriately we need to first formalize the processes and dynamics through which algorithmic decisions impact their subjects and society in the long run.

1.1. Related Work

Most existing notions of algorithmic fairness are one-shot and require that a particular error metric is equal across all social groups. Different choices for the metric have led to different fairness criteria; examples include demographic parity (Kleinberg et al., 2017; Dwork et al., 2012; Corbett-Davies et al., 2017), disparate impact (Zafar et al., 2017; Feldman et al., 2015), equality of odds (Hardt et al., 2016), and calibration (Kleinberg et al., 2017). Prior notions fail to capture the disparity in the *effort* it takes members of different social groups to improve their algorithmic outcome. We propose a group-dependent, data-driven measure of effort, inspired by the literature on Equality Of Opportunity (EOP) (Roemer & Trannoy, 2015; Heidari et al., 2019). (In particular, the effort it takes individual i to improve their feature k value from x to x' is proportional to the difference between the rank/quantile of x' and x in the distribution of feature k in i ’s social group.)

Social (or observational) learning (Bandura, 2008) is a type of learning that occurs through observing and imitating the behavior of others. This type of learning requires a social model (or role model)—someone of higher status in the environment. According to the social learning theory, observers recreate their role model’s behavior only if they have sufficient motivation (Bandura, 1962; Apesteguia et al., 2007)—this often comes from the observation that the model is rewarded for their actions. In our model, an individual recreates their role model’s qualification if by doing so he obtains a positive utility, where utility is defined as reward minus effort. Furthermore, it has been shown that observers learn best from models that they identify with² and find it within their capability to imitate them (Bandura, 1962). These points are captured by our effort function. Social learning explicitly captures the role model implications of decision making policies. This echoes research in sociology and economics which has already established the role model effects of affirmative action policies (Chung, 2000). We note that imitation dynamics have been extensively studied in population and evolutionary games (see, e.g., (Sandholm, 2010), Chapters 4 and 5).

Several recent papers study the impact of decision-making models and fairness interventions on society and individuals (see, e.g., (Liu et al., 2018; Kannan et al., 2019)). Unlike prior work, our focus is on *how subjects respond* to algorithmic policies by *improving/updating their (mutable) qual-*

²Social identity is a person’s sense of who they are based on their group membership(s) (Tajfel et al., 1979).

ifications. We don't make any case-specific assumptions about how the world changes in response to the deployed model, rather allow our micro-scale behavioral model to derive the macro-level change. We emphasize that our model is not meant to perfectly capture all the behavioral nuances involved; rather our primary goal is to highlight the potential role of behavioral dynamics and human responses in shaping the long-term impact of algorithmic models.

Also related but orthogonal to our work is a recent line of research on *strategic classification*—a setting in which decision subjects are assumed to respond *strategically* and potentially *untruthfully* to the choice of the classification model, and the goal is to design classifiers that are robust to strategic manipulation (Dong et al., 2018; Hu et al., 2019; Milli et al., 2019).

2. Setting

We consider the standard supervised learning setting: A learning algorithm receives the training data set $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ consisting of n instances, where $\mathbf{x}_i \in \mathcal{X}$ specifies the feature vector for individual i and $y_i \in \mathcal{Y}$, the *ground truth* label for him/her. We use s_i to refer to the sensitive feature value (e.g., race, gender, or their intersection) for individual i . For the ease of notation, we will use \mathbf{z}_i to denote the example (\mathbf{x}_i, y_i) . We assume \mathbf{z}_i fully characterizes individual i with respect to the task at hand. The training data is sampled i.i.d. from a distribution P on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. For simplicity, throughout we assume there exists an unknown function $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that for all i , $y_i = f(\mathbf{x}_i)$. Unless specified otherwise, we assume $\mathcal{X} \subseteq \mathbb{R}^K$, where K denotes the number of features. The goal of a learning algorithm is to use the training data to fit a (regression) *model* (or hypothesis) $h : \mathcal{X} \rightarrow \mathcal{Y}$ that accurately predicts the label for new instances. Let \mathcal{H} be the hypothesis class consisting of all the models available to the learning algorithm. A learning algorithm receives D as the input; then utilizes the data to select a model $h \in \mathcal{H}$ that minimizes some notion of loss, L . For instance, in regression the empirical mean squared loss of a model h on D is defined as $L_D(h) = \sum_{i \in D} (y_i - \hat{y}_i)^2$, where $\hat{y}_i = h(\mathbf{x}_i)$. The learning algorithm outputs the model $h \in \mathcal{H}$ that minimizes the empirical loss; i.e., $h = \arg \min_{h'} L_D(h')$.

We assume there exists a benefit function $b : \mathcal{X}, \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ that quantifies the benefit an individual with feature vector \mathbf{x} and ground truth label y receives if the trained model predicts label \hat{y} for them. Throughout this work we will focus on benefit functions that are only functions of y, \hat{y} and are linear in \hat{y} (e.g., $b(y, \hat{y}) = \hat{y} - y + 1$ or $b(y, \hat{y}) = \hat{y}$). For simplicity and ease of interpretation, all illustrations in the main body of the paper are performed with \hat{y} as the benefit function. Throughout, we assume higher predicted labels are considered more desirable from the point of view of

individual decision subjects (e.g., this is the case when the task is to predict students' grade to decide who is admitted to a top school).

2.1. Efforts, Rewards, Utilities

Let h specify the deployed predictive model. Consider an individual characterized by $\mathbf{z} = (\mathbf{x}, y)$. Let $\mathcal{R}_h(\mathbf{z}, \mathbf{z}')$ specify the *reward* or added benefit he/she obtains as the result of changing his/her characteristics from \mathbf{z} to \mathbf{z}' :

$$\mathcal{R}_h(\mathbf{z}, \mathbf{z}') = b(h(\mathbf{x}'), y')^\alpha - b(h(\mathbf{x}), y)^\alpha,$$

where $\alpha > 0$ is a constant specifying the individual's degree of risk aversion. This parameter can be adjusted to model diminishing returns to added benefit. Unless otherwise specified, in our illustrations we take $\alpha = 1$. Let $\mathcal{E}_h(\mathbf{z}, \mathbf{z}')$ specify the effort it takes the individual to update their qualifications and make the change from \mathbf{z} to \mathbf{z}' (we will shortly elaborate on how effort can be quantified). The overall utility of the individual is denoted by $\mathcal{U}_h(\mathbf{z}, \mathbf{z}')$ and for simplicity, we assume it takes on a linear form:

$$\mathcal{U}_h(\mathbf{z}, \mathbf{z}') = \mathcal{R}_h(\mathbf{z}, \mathbf{z}') - \mathcal{E}_h(\mathbf{z}, \mathbf{z}').$$

That is, utility is simply reward minus effort. When clear from the context, we drop the subscript h .

Throughout, we focus on effort functions that only depend on \mathbf{x}, \mathbf{x}' , s . For simplicity, we assume the effort function is additively separable across features—that is, the total effort required to change \mathbf{x} to \mathbf{x}' is a linear combination of the effort needed to change each feature separately:³

$$\mathcal{E}(\mathbf{z}, \mathbf{z}') = \mathcal{E}_s(\mathbf{x}, \mathbf{x}') = c_s + \frac{1}{K} \sum_{k=1}^K c_{s,k} \epsilon_{s,k}(x_k, x'_k),$$

where $\epsilon_{s,k}(x_k, x'_k)$ denotes the effort it takes to change the value of feature k from x_k to x'_k for an individual belonging to group s (defined below). For group s , c_s is a group-dependent constant specifying the minimum effort required to make any change. Similarly for $k = 1, \dots, K$, $c_{s,k}$'s are constant weights which allow us to specify the relative difficulty of change across different features. For simplicity, throughout we assume $c_{s,k} \equiv 1$ for all k, s .

We define $\epsilon_{s,k}(x_k, x'_k)$ as follows—depending on the feature type and our domain knowledge about the feature:⁴

³Depending on our domain knowledge of how features relate to one another, we may find a different aggregating operator (e.g., max) to be more appropriate than summation. For instance, if two features automatically change together, no extra effort is required for changing both of them simultaneously. Throughout our illustrations, for simplicity we focus on additively separable functions, but our results can be readily produced for more complicated effort functions.

⁴See the Appendix for a complete description of the effort function.

Suppose feature k is numerical and *monotone* (Duivesteijn & Feelders, 2008), that is, we expect an increase in its value to monotonically increase the predicted label—everything else being equal. Monotonicity implies that there is a clear direction of change that is considered desirable. As an example in the education context, consider the *number of hours of study*: we expect an increase in this feature to increase the student’s predicted grade. Without loss of generality, we assume higher values of feature k are expected to increase the predicted label (we can ensure this by preprocessing the data and negating feature k values if necessary). We define $\epsilon_{s,k}(x_k, x'_k)$ as follows:

$$\epsilon_{s,k}(x_k, x'_k) = \max\{0, Q_{s,k}(x'_k) - Q_{s,k}(x_k)\}.$$

Above, $Q_{s,k}(x)$ specifies the quantile/rank of value x in the empirical distribution of feature k among individuals who belong to group s .

The above effort function is inspired by the line of work on equality of opportunity (Roemer & Trannoy, 2015): Note that the distribution of feature k values can be different across different social groups (e.g., men and women, or African-Americans and Whites). We take the view that this is potentially *not* the result of one group being inherently inferior to another (in terms of feature k). Rather, it is most likely due to the underlying socio-economic circumstances that the privileged group can achieve higher values of feature k with less effort. To account for this in our effort function, we measure the effort it takes a person i in group s to change their feature k value from x_k to x'_k by comparing the rank/quantile of x'_k and x_k in the empirical distribution of feature k within group s . This implies that if for most people in group s , the value of feature k is equal or better than x'_k , we consider it relatively easy for individual i to make the change from x_k to x'_k . If however, very few in group s have ever been able to achieve x'_k , then it is considered very difficult for i to make this change.

Using the effort function defined above, in Section 3 we propose a new effort-based measure of algorithmic unfairness. In Section 4, we will utilize our effort function to compute individual utilities and subsequently specifying the social models.

3. Effort-based Measures of Fairness

Existing formulations of fairness are concerned with how errors are distributed among various social groups, but they do not account for the fact that even if errors are distributed similarly, the *effort* required to fix those errors and improve one’s prediction may be significantly higher for the disadvantaged subpopulation. Building on our notion of effort introduced in Section 2, in this section we formulate a new measure of algorithmic unfairness, called the *effort unfairness*. At a high level, effort unfairness is the disparity in

the average effort members of each group have to exert to obtain their desired outcomes—by imitating the appropriate role models.⁵

We propose three different formalizations of effort unfairness: *bounded-effort* unfairness, *threshold-reward* unfairness, and *effort-reward* unfairness. Each notion corresponds to a distinct/salient way in which decision subjects may evaluate fairness and respond to their predictions.




Figure 1: The average reward obtainable by members of each group as the result of exerting (at most) δ units of effort. For each model, the bounded effort unfairness is the difference between the dashed and solid curves.

Bounded-effort unfairness is the inequality in the average reward members of each group can obtain by exerting a fixed level of effort. More precisely:

Definition 1 (Bounded-effort Unfairness) *Given a constant δ , the δ -bounded-effort unfairness of a predictive model h is the inequality of the following metric across different groups:*

$$\mathbb{E}_{i \sim P: s_i = s} \left(\max_{\mathbf{z} \in \mathcal{Z}} \mathcal{R}_h(\mathbf{z}_i, \mathbf{z}) \text{ s.t. } \mathcal{E}_h(\mathbf{z}_i, \mathbf{z}) \leq \delta \right). \quad (1)$$

The bounded-effort formulation is motivated by the literature on *bounded willpower* in behavioral economics (Mullainathan & Thaler, 2000), which at a high level posits that there is an upper bound on the level of effort people can be expected to exert.

To compute the bounded-effort unfairness in practice, we propose replacing the expectation in Equation 1 with the empirical mean, and taking the maximum over the available data set D . The latter not only simplifies the optimization,

⁵Note that algorithmic unfairness has many different aspects. We introduce a new dimension along which algorithmic decisions disparately affect different subpopulation, but this is not to undermine the importance of all other dimensions of unfairness (such as error disparities). In particular, we do *not* claim that if a model is fair according to our notion, then it cannot be unfair according to other criteria.

it also has a natural interpretation in terms of *social models* (see Section 4): Precisely, we estimate Equation 1 by $\frac{1}{n_s} \sum_{i:s_i=s} (\max_{\mathbf{z} \in D} \mathcal{R}_h(\mathbf{z}_i, \mathbf{z}) \text{ s.t. } \mathcal{E}_h(\mathbf{z}_i, \mathbf{z}) \leq \delta)$ where n_s is the number of subjects in group s .

Figure 1 illustrates the bounded effort unfairness calculated over the student performance data set (Cortez & Silva, 2008). (Information about the data set, our pre-processing steps, and the trained models can be found in the Appendix. The code used to generate the plots in this paper can be found at https://github.com/nvedant07/effort_reward_fairness.) Note that according to the bounded-effort measure, different models may be discriminatory against different groups. Also depending on the choice of δ the measure may rank the three models differently.

Threshold-reward unfairness is the inequality in the average effort member of each group need to exert to reach a certain level of reward. More precisely:

Definition 2 (Threshold-reward Unfairness) Given a constant δ , the δ -threshold-reward unfairness of a predictive model h is the inequality of the following metric across different groups:

$$\mathbb{E}_{i \sim P:s_i=s} \left(\min_{\mathbf{z} \in \mathcal{Z}} \mathcal{E}_h(\mathbf{z}_i, \mathbf{z}) \text{ s.t. } \mathcal{R}_h(\mathbf{z}_i, \mathbf{z}) \geq \delta \right). \quad (2)$$

This formulation is motivated by the *capability* view of fairness (Sen, 1993): Sen conceptualizes fairness as the equality of capability, where at a high level, capability is a person’s ability to reach valuable states of being (in our case, a certain level of reward).

Figure 2 illustrates the average effort unfairness on the student performance data set. Note that depending on the choice of δ the measure may rank the three models differently. Also interestingly, depending on the choice of δ the same model (i.e., linear model) may be considered unfair toward men, unfair toward women, or perfectly fair!

The previous two formulations of effort unfairness—while well-motivated—may give us different rankings across the same set of alternatives depending on our choice of δ . The final formulation, which we call the *effort-reward* unfairness, resolves this issue by comparing the highest utility members of each group can possibly achieve by exerting additional effort. More precisely:

Definition 3 (Effort-reward Unfairness) For a predictive model h , the *effort-reward unfairness* is the inequality of the following metric across different groups:

$$\mathbb{E}_{i \sim P:s_i=s} \max_{\mathbf{z} \in \mathcal{Z}} \mathcal{U}_h(\mathbf{z}_i, \mathbf{z}). \quad (3)$$

Figure 3 contrasts the effort-reward measure with existing




Figure 2: The average effort required from members of each group to obtain (at least) δ units of reward. For each model, the threshold-reward unfairness is the difference between the dashed and solid curves.

notions of algorithmic unfairness (these measures are precisely defined in the Appendix.) As evident in Figure 3, on the student performance data set none of the existing fairness notions fully captures the effort disparity. See the Appendix for a numerical example further illustrating why and how our measure of effort unfairness may not be captured by existing notions.




Figure 3: Comparison of the effort-reward unfairness with several existing notions of (un)fairness. Effort-reward unfairness ranks the models differently.

4. How Algorithmic Policies Re-shape Society

To formulate the long-term impact of algorithmic policies on the underlying population, in this section we propose a micro-scale model of how individuals may respond to algorithmic policies, taking inspiration from the psychological literature on *social learning* (Bandura, 1978). We posit that individuals observe and potentially imitate the behavior of their so-called *social models*. A social model is another decision subject who has received a higher level of benefit as the result of being subject to the decision-making model. Our social learning model captures settings in which subjects don’t know the inner workings of the decision making




Figure 4: Segregation measures computed over the initial and impacted population for 3 models. Segregation may increase or decrease as the result of the behavioral dynamics.

model, but can infer how to improve their standing by observing the decisions it makes for people similar to them (i.e., their social models).

With the behavioral dynamics specified, we can quantify the macro-scale long-term impact of the model on reshaping the underlying populations. We adapt existing measures of segregation from sociology and economics (Massey & Denton, 1988) to characterize how the distribution of qualifications for each group changes in response to the deployed model. Measures of segregation quantify how separate the two subpopulations are in terms of distribution of qualifications. We believe such *model-independent* measures are important to consider, because the decision making model itself may change over time, but its impact on the underlying population may be long lasting.

4.1. Behavioral Dynamics

At a high level, we simulate every individual’s response to the predictive model by selecting a social model for them from the training data set; the individual is then assumed to exert effort to attain his/her model’s qualifications if and only if doing so improves her overall utility.

Our micro-scale model is meant to capture two important nuances pointed out by the social learning theory: First, according to the theory observers recreate their social model’s behavior only if they have sufficient motivation and this motivation often comes from observing that the social model is rewarded for their actions. We capture this by assuming that an individual recreates his/her social model’s qualification if by doing so he/she is sufficiently rewarded and obtains a positive utility. Second, it has been shown that observers learn best from social models that they identify with and find it within their ability to emulate. These points are captured through our notion of *effort* and *utility*. If a potential social model belongs to a different group than that of the individual, the effort it takes to recreate his/her actions is very high, therefore the individual won’t find sufficient utility in imitating him/her.

Assuming that the training data is a representative sample of the population⁶, we select the social models from among the individuals present in the data set. In particular, for an individual \mathbf{z} , the model (denoted by \mathbf{z}') is another decision subject in D whose imitation would maximize \mathbf{z} ’s utility. That is, $\mathbf{z}' = \arg \max_{\mathbf{z}'' \in D} \mathcal{U}_h(\mathbf{z}, \mathbf{z}'')$.

Two remarks are in order. First, note that each one of our fairness notions corresponds to a criterion for choosing the social model. Depending on the context, one criterion may better reflect the human response. In this paper, we deliberately focus on utility maximization. This choice is primarily for ease of illustration—it allows us to forgo specifying δ —but our analysis can be replicated for the other two criteria, as well.

Second, one may ask how can individuals be expected to find the right social model (in particular, the utility-maximizing one)? One concrete way is through *actionable* or *counterfactual* explanations (Wachter et al., 2017; Ustun et al., 2018). When an individual fails to receive their desired prediction, such explanations lay out the optimal change he/she can make to improve their outcome. (We re-emphasize that we do not consider our model to be a perfect reflection of extremely nuanced human behavior in the real world. We do, however, consider it to be a reasonable approximation of certain aspects of the process. Our primary goal with this model is to illustrate the role of behavioral dynamics in driving the societal impact of a decision making model.)

Through our proposed dynamics, we can simulate how subjects respond to the predictive model. We then obtain a new data set D' representing the *impacted population*’s qualifications. Next, we adopt measures of segregation to compare the initial and impacted populations and quantify the long-term macro-scale impact of the predictive model on the underlying population. Our interest in measures of segregation is rooted in the observation that unfairness usually emerges as a concern when there is a clear separation

⁶We cannot always make this assumption—the sampling process can become biased in numerous ways.

between different segments of the population (in terms of qualifications and/or outcomes). If people belonging to two socially salient groups are fully mixed—in terms of their features and outcomes—such concerns are unlikely to arise. We emphasize that segregation does not always imply unfairness (for instance, segregation could be the consequence of *specialization*: different segments of the population may willingly invest in different sets of qualifications). But unfairness often comes with some form of segregation. We, therefore, propose measures of segregation as an effective test for potential unfairness.

4.2. Quantifying Impact via Measures of Segregation

Ethnic and racial segregation is a well-studied phenomenon in sociology. At a high level, segregation is the degree to which two or more groups live separately from one another. A long line of work in sociology has been concerned with measuring segregation. In their highly influential article, [Massey & Denton \(1988\)](#) break down residential segregation into five distinct axes of measurement: centralization, evenness, clustering, exposure, and concentration. Below, we overview these measures and show how three of them can be utilized to measure the macro-scale impact of decision-making models on the distribution of qualifications.

Evenness measures how unevenly the minority group is distributed over areal units. Evenness is maximized when all units have the same relative number of minority and majority members as the whole population. More precisely, for an area/neighborhood i , let t_i denote its total population, m_i the number of minority residents, and M_i the number of majority residents of the neighborhood. Also, let $p_i = m_i/t_i$ specify the percentage of minority residents in the area. Let T and P specify the total population size and minority proportion of the whole population. Suppose there are N areal units in total. The Atkinson Index (AI) is a particular measure of evenness satisfying several desirable properties.⁷ For a constant $0 < \beta < 1$, the Atkinson index measures the inequality of $(1 - p_i)/p_i$ (i.e. the number of majority residents per minority resident in neighborhood i) computed across all individuals belonging to the minority group: $1 - \frac{P}{1-P} \left(\frac{1}{N} \sum_{i=1}^N (1 - p_i)^{1-\beta} p_i^\beta t_i / TP \right)^{1/(1-\beta)}$.

Centralization is the degree to which a group is spatially located near the center of an urban area. (Because of certain urban development policies in the past, central areas of most cities across the U.S. are declining residential areas.) The degree of centralization can be measured by comparing the percentage of minority residents living in the central areas of the city. The Centralization Index (CI) is precisely defined as follows: $\frac{\sum_{i \text{ central}} m_i}{m}$ where m is the total minority

⁷It satisfies the *transfer principle*, *compositional invariance*, *population invariance*, and *organizational equivalence*.

population.

Clustering measures the extent to which areal units inhabited by minority members adjoin one another, or cluster, in space. For example, the *Absolute Clustering Index* (ACI) “expresses the average number of [minority] members in nearby [areal units] as a proportion of the total population in those nearby [areal units]” ([Massey & Denton, 1988](#)). ACI is defined as follows:

$$\frac{\left[\sum_{i=1}^N \sum_{j=1}^N c_{i,j} \frac{m_i m_j}{m} \right] - \left[\sum_{i=1}^N \sum_{j=1}^N c_{i,j} \frac{1}{N} \frac{m}{N} \right]}{\left[\sum_{i=1}^N \sum_{j=1}^N c_{i,j} \frac{m_i t_j}{m} \right] - \left[\sum_{i=1}^N \sum_{j=1}^N c_{i,j} \frac{1}{N} \frac{m}{N} \right]}.$$

For any two areas i and j , $c_{i,j}$ specifies the closeness between their corresponding centers.

Residential **exposure** refers to the degree of potential contact, or the possibility of interaction, between minority and majority group members within geographic areas of a city. **Concentration** refers to relative amount of physical space occupied by a minority group in the urban environment.

There are several notable differences between our setting and that of residential segregation. First, in our setting there are no predefined notions of “area” or “neighborhoods” that individuals belong to. Second, individuals are described by multi-dimensional feature vectors—as opposed to a 2-dimensional vector specifying their residential location. Third, it is not immediately clear how distances and similarity between individuals should be defined. Next, we will address these issues for evenness, centralization, and clustering.

Throughout this section, we will focus on the mutable feature subspace. We take the distance between two individuals i (belonging to group s) and j (belonging to group s') as follows: $d(i, j) = \max\{d^s(i, j), d^{s'}(i, j)\}$, where $d^s(i, j)$ is defined as follows: $\max\{0, Q_s(y_j) - Q_s(y_i)\} + \sum_{k \in \text{mutable}} \epsilon_{s,k}(x_k, x'_k)$. We will use Atkinson index to measure evenness. We will specify areas through what we call *focal points*—these are feature vectors in the mutable feature subspace that at least one subject in the original population imitates to improve their utility. Each focal point corresponds to a neighborhood, and an individual belongs to the neighborhood of their nearest focal point. We measure the degree of centralization by comparing the percentage of minority individuals whose predictions are above the average (e.g., $\hat{Y} > 11.94$). We measure ACI at the individual level—that is, we assume each individual corresponds to a neighborhood. We define the similarity between two neighborhoods as follows: $c_{i,j} = e^{-d(i,j)}$.

Figure 4 illustrates our measures of segregation both for the initial population (depicted in blue) and the impacted population—after individuals respond to the model by adjusting their qualifications (depicted in red). Segregation




Figure 5: The change in various measures of segregation as a function of enforcing fairness constraints with strength τ .

can change in counter-intuitive ways through imitation dynamics.

5. The Role of Fairness Interventions

In this Section, we investigate the effect of enforcing fairness constraints—at the time of training—on the long term population-level impact of the deployed model. We focus on the case of linear regression. We train a model by minimizing the mean squared error while imposing the welfare constraints proposed by Heidari et al. (2018).

Figure 5 shows the effect of imposing fairness constraints on various measures of segregation (all computed by taking females as the minority/protected group). One might expect that these constraints would always reduce segregation in the long run. As illustrated in Figure 5, this is not always the case. For a small value of τ , enforcing fairness constraints can significantly reduce the degree of clustering (see Figure 5a). Larger values of τ can reverse this effect and lead to a population that is more heavily clustered/segregated compared to the original population. Evenness remains relatively unchanged regardless of the value of τ (see Figure 5b).

These findings highlight an important insight about fairness constraints: they can affect segregation in two competing ways. On the one hand by automatically assigning a desirable label to some members of the disadvantaged group, the model incentivizes these members not to make any change. On the other hand, these members can serve as social models for the rest of the disadvantaged group, nudging more of them to improve their qualifications and obtain better labels. Which force is more powerful? One can only answer this by simulating the dynamics on the particular data set at hand. We see clear parallels between our observations and the prior work on affirmative action policies. Advocates of affirmative action often argue that a larger representation of minorities in desirable positions can lead to role models who encourage other minorities in their investment decisions (see e.g., (Chung, 2000)). At the same time, critics

argue that affirmative action quota may indirectly harm the disadvantaged group members by reducing their incentives to invest in qualifications (Coate & Loury, 1993a;b). Similar to our work, economic results on the long-term impact of affirmative action policies is mixed and context-specific.

We end this section with a remark on fairness-restoring interventions. While we focused on *algorithmic* interventions, we must emphasize that changing the decision-making model is not the only mechanism through which segregation and unfairness can be alleviated. Instead of artificially changing the decision boundary, it may be socially more desirable to address unfairness *before* people are subjected to algorithmic decision making. For instance, one could design and implement policies that make it easier for disadvantaged group members to obtain certain qualifications. We leave the analysis of such *feature interventions* as a promising direction for future work.

6. Conclusion & Future Directions

We presented a data-driven framework for studying the potential long-term impact of predictive models on decision subjects and society. We proposed a micro-model of human response to algorithmic policies rooted in psychology and several macro-level measures of change borrowed from sociology and economics. Our work suggests several immediate directions for future work, including but not limited to (a) human subject experiments to investigate the viability of our behavioral model; (b) designing an efficient mechanism for bounding effort-reward unfairness.

Acknowledgements

K. P. Gummadi is supported in part by an ERC Advanced Grant “Foundations for Fair Social Computing” (no. 789373).

References

- Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias. *ProPublica*, 2016.
- Apestequia, J., Huck, S., and Oechssler, J. Imitation— theory and experimental evidence. *Journal of Economic Theory*, 136(1):217–235, 2007.
- Bandura, A. Social learning through imitation. 1962.
- Bandura, A. Social learning theory of aggression. *Journal of communication*, 28(3):12–29, 1978.
- Bandura, A. Observational learning. *The international encyclopedia of communication*, 2008.
- Barry-Jester, A., Casselman, B., and Goldstein, D. The new science of sentencing. *The Marshall Project*, August 2015.
- Chung, K.-S. Role models and arguments for affirmative action. *American Economic Review*, 90(3):640–648, 2000.
- Coate, S. and Loury, G. Antidiscrimination enforcement and the problem of patronization. *The American Economic Review*, 83(2):92–98, 1993a.
- Coate, S. and Loury, G. C. Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, pp. 1220–1240, 1993b.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 797–806. ACM, 2017.
- Cortez, P. and Silva, A. M. G. Using data mining to predict secondary school student performance. 2008.
- Deo, R. C. Machine learning in medicine. *Circulation*, 132(20):1920–1930, 2015.
- Dong, J., Roth, A., Schutzman, Z., Waggoner, B., and Wu, Z. S. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 55–70. ACM, 2018.
- Duivesteyn, W. and Feelders, A. Nearest neighbour classification with monotonicity constraints. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 301–316. Springer, 2008.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the Innovations in Theoretical Computer Science Conference*, pp. 214–226. ACM, 2012.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, pp. 259–268. ACM, 2015.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Proceedings of the 30th Conference on Neural Information Processing Systems*, pp. 3315–3323, 2016.
- Heidari, H., Ferrari, C., Gummadi, K. P., and Krause, A. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Proceedings of the 32nd Conference on Neural Information Processing Systems*, 2018.
- Heidari, H., Loi, M., Gummadi, K. P., and Krause, A. A moral framework for understanding of fair ml through economic models of equality of opportunity. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, 2019.
- Hu, L., Immorlica, N., and Vaughan, J. W. The disparate effects of strategic manipulation. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, 2019.
- Kannan, S., Roth, A., and Ziani, J. Downstream effects of affirmative action. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, 2019.
- Kleinberg, J., Mullainathan, S., and Raghavan, M. Inherent trade-offs in the fair determination of risk scores. In *In proceedings of the 8th Innovations in Theoretical Computer Science Conference*, 2017.
- Levin, S. A beauty contest was judged by AI and the robots didn’t like dark skin. *The Guardian*, 2016.
- Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., and Hardt, M. Delayed impact of fair machine learning. In *Proceedings of the International Conference on Machine Learning*, 2018.
- Massey, D. S. and Denton, N. A. The dimensions of residential segregation. *Social forces*, 67(2):281–315, 1988.
- Milli, S., Miller, J., Dragan, A. D., and Hardt, M. The social cost of strategic classification. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, 2019.
- Mullainathan, S. and Thaler, R. H. Behavioral economics. Technical report, National Bureau of Economic Research, 2000.

- Petrasic, K., Saul, B., Greig, J., and Bornfreund, M. Algorithms and bias: What lenders need to know. *White & Case*, 2017.
- Roemer, J. E. and Trannoy, A. Equality of opportunity. In *Handbook of income distribution*, volume 2, pp. 217–300. Elsevier, 2015.
- Rudin, C. Predictive policing using machine learning to detect patterns of crime. *Wired Magazine*, August 2013. Retrieved 4/28/2016.
- Sandholm, W. H. *Population games and evolutionary dynamics*. MIT press, 2010.
- Sen, A. Capability and well-being. *The quality of life*, 30, 1993.
- Sweeney, L. Discrimination in online ad delivery. *Queue*, 11(3):10, 2013.
- Tajfel, H., Turner, J. C., Austin, W. G., and Worchel, S. An integrative theory of intergroup conflict. *Organizational identity: A reader*, pp. 56–65, 1979.
- Ustun, B., Spangher, A., and Liu, Y. Actionable recourse in linear classification. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, 2018.
- Wachter, S., Mittelstadt, B., and Russell, C. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. 2017.
- Zafar, M. B., Valera, I., Gomez Rodriguez, M., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.