

Delayed Impact of Fair Machine Learning

Lydia T. Liu* Sarah Dean* Esther Rolf* Max Simchowitz* Moritz Hardt*

April 10, 2018

Abstract

Fairness in machine learning has predominantly been studied in static classification settings without concern for how decisions change the underlying population over time. Conventional wisdom suggests that fairness criteria promote the long-term well-being of those groups they aim to protect.

We study how static fairness criteria interact with temporal indicators of well-being, such as long-term improvement, stagnation, and decline in a variable of interest. We demonstrate that even in a one-step feedback model, common fairness criteria in general do not promote improvement over time, and may in fact cause harm in cases where an unconstrained objective would not. We completely characterize the delayed impact of three standard criteria, contrasting the regimes in which these exhibit qualitatively different behavior. In addition, we find that a natural form of measurement error broadens the regime in which fairness criteria perform favorably.

Our results highlight the importance of measurement and temporal modeling in the evaluation of fairness criteria, suggesting a range of new challenges and trade-offs.

1 Introduction

Machine learning commonly considers static objectives defined on a snapshot of the population at one instant in time; consequential decisions, in contrast, reshape the population over time. Lending practices, for example, can shift the distribution of debt and wealth in the population. Job advertisements allocate opportunity. School admissions shape the level of education in a community.

Existing scholarship on fairness in automated decision-making criticizes unconstrained machine learning for its potential to *harm* historically underrepresented or disadvantaged groups in the population [Executive Office of the President, 2016, Barocas and Selbst, 2016]. Consequently, a variety of *fairness criteria* have been proposed as constraints on standard learning objectives. Even though, in each case, these constraints are clearly intended to *protect* the disadvantaged group by an appeal to intuition, a rigorous argument to that effect is often lacking.

In this work, we formally examine under what circumstances fairness criteria do indeed promote the long-term well-being of disadvantaged groups measured in terms of a temporal variable of interest. Going beyond the standard classification setting, we introduce a one-step feedback model of decision-making that exposes how decisions change the underlying population over time.

*Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

Our running example is a hypothetical lending scenario. There are two groups in the population with features described by a summary statistic, such as a *credit score*, whose distribution differs between the two groups. The bank can choose thresholds for each group at which loans are offered. While group-dependent thresholds may face legal challenges [Ross and Yinger, 2006], they are generally inevitable for some of the criteria we examine. The impact of a lending decision has multiple facets. A default event not only diminishes profit for the bank, it also worsens the financial situation of the borrower as reflected in a subsequent decline in credit score. A successful lending outcome leads to profit for the bank and also to an increase in credit score for the borrower.

When thinking of one of the two groups as disadvantaged, it makes sense to ask what lending policies (choices of thresholds) lead to an expected improvement in the score distribution within that group. An unconstrained bank would maximize profit, choosing thresholds that meet a break-even point above which it is profitable to give out loans. One frequently proposed fairness criterion, sometimes called demographic parity, requires the bank to lend to both groups at an equal rate. Subject to this requirement the bank would continue to maximize profit to the extent possible. Another criterion, originally called equality of opportunity, equalizes the *true positive rates* between the two groups, thus requiring the bank to lend in both groups at an equal rate among individuals who repay their loan. Other criteria are natural, but for clarity we restrict our attention to these three.

Do these fairness criteria benefit the disadvantaged group? When do they show a clear advantage over unconstrained classification? Under what circumstances does profit maximization work in the interest of the individual? These are important questions that we begin to address in this work.

1.1 Contributions

We introduce a one-step feedback model that allows us to quantify the long-term impact of classification on different groups in the population. We represent each of the two groups A and B by a *score* distribution π_A and π_B , respectively. The support of these distributions is a finite set $\mathcal{X}$ corresponding to the possible values that the score can assume. We think of the score as highlighting one variable of interest in a specific domain such that higher score values correspond to a higher probability of a positive outcome. An *institution* chooses selection policies $\tau_A, \tau_B: \mathcal{X} \rightarrow [0, 1]$ that assign to each value in $\mathcal{X}$ a number representing the rate of selection for that value. In our example, these policies specify the lending rate at a given credit score within a given group. The institution will always maximize their utility (defined formally later) subject to either (a) no constraint, or (b) equality of selection rates, or (c) equality of true positive rates.

We assume the availability of a function $\Delta: \mathcal{X} \rightarrow \mathbb{R}$ such that $\Delta(x)$ provides the expected change in score for a selected individual at score x . The central quantity we study is the expected difference in the mean score in group $j \in \{A, B\}$ that results from an institutions policy, $\Delta\mu_j$ defined formally in Equation (2). When modeling the problem, the expected mean difference can also absorb external factors such as “reversion to the mean” so long as they are mean-preserving. Qualitatively, we distinguish between *long-term improvement* ($\Delta\mu_j > 0$), *stagnation* ($\Delta\mu_j = 0$), and *decline* ($\Delta\mu_j < 0$). Our findings can be summarized as follows:

1. Both fairness criteria (equal selection rates, equal true positive rates) can lead to all possible outcomes (improvement, stagnation, and decline) in natural parameter regimes. We provide a complete characterization of when each criterion leads to each outcome in Section 3.

- There are a class of settings where equal selection rates cause decline, whereas equal true positive rates do not (Corollary 3.5),
 - Under a mild assumption, the institution’s optimal unconstrained selection policy can never lead to decline (Proposition 3.1).
2. We introduce the notion of an *outcome curve* (Figure 1) which succinctly describes the different regimes in which one criterion is preferable over the others.
 3. We perform experiments on FICO credit score data from 2003 and show that under various models of bank utility and score change, the outcomes of applying fairness criteria are in line with our theoretical predictions.
 4. We discuss how certain types of measurement error (e.g., the bank underestimating the repayment ability of the disadvantaged group) affect our comparison. We find that measurement error narrows the regime in which fairness criteria cause decline, suggesting that measurement should be a factor when motivating these criteria.
 5. We consider alternatives to hard fairness constraints.
 - We evaluate the optimization problem where fairness criterion is a regularization term in the objective. Qualitatively, this leads to the same findings.
 - We discuss the possibility of optimizing for group score improvement $\Delta\mu_j$ directly subject to institution utility constraints. The resulting solution provides an interesting possible alternative to existing fairness criteria.

We focus on the impact of a selection policy over a single epoch. The motivation is that the designer of a system usually has an understanding of the time horizon after which the system is evaluated and possibly redesigned. Formally, nothing prevents us from repeatedly applying our model and tracing changes over multiple epochs. In reality, however, it is plausible that over greater time periods, economic background variables might dominate the effect of selection.

Reflecting on our findings, we argue that careful temporal modeling is necessary in order to accurately evaluate the impact of different fairness criteria on the population. Moreover, an understanding of measurement error is important in assessing the advantages of fairness criteria relative to unconstrained selection. Finally, the nuances of our characterization underline how intuition may be a poor guide in judging the long-term impact of fairness constraints.

1.2 Related work

Recent work by Hu and Chen [2018] considers a model for long-term outcomes and fairness in the labor market. They propose imposing the demographic parity constraint in a *temporary* labor market in order to provably achieve an equitable long-term equilibrium in the *permanent* labor market, reminiscent of economic arguments for affirmative action [Foster and Vohra, 1992]. The equilibrium analysis of the labor market dynamics model allows for specific conclusions relating fairness criteria to long term outcomes. Our general framework is complementary to this type of domain specific approach.

Fuster et al. [2017] consider the problem of fairness in credit markets from a different perspective. Their goal is to study the effect of machine learning on interest rates in different groups at an equilibrium, under a static model without feedback.

Ensign et al. [2017] consider feedback loops in predictive policing, where the police more heavily monitor high crime neighborhoods, thus further increasing the measured number of crimes in those neighborhoods. While the work addresses an important temporal phenomenon using the theory of urns, it is rather different from our one-step feedback model both conceptually and technically.

Demographic parity and its related formulations have been considered in numerous papers [e.g. Calders et al., 2009, Zafar et al., 2017]. Hardt et al. [2016] introduced the equality of opportunity constraint that we consider and demonstrated limitations of a broad class of criteria. Kleinberg et al. [2017] and Chouldechova [2016] point out the tension between “calibration by group” and equal true/false positive rates. These trade-offs carry over to some extent to the case where we only equalize true positive rates [Pleiss et al., 2017].

A growing literature on fairness in the “bandits” setting of learning [see Joseph et al., 2016, *et sequelae*] deals with online decision making that ought not to be confused with our one-step feedback setting. Finally, there has been much work in the social sciences on analyzing the effect of affirmative action [see e.g., Keith et al., 1985, Kaley et al., 2006].

1.3 Discussion

In this paper, we advocate for a view toward long-term outcomes in the discussion of “fair” machine learning. We argue that without a careful model of delayed outcomes, we cannot foresee the impact a fairness criterion would have if enforced as a constraint on a classification system. However, if such an accurate outcome model is available, we show that there are more direct ways to optimize for positive outcomes than via existing fairness criteria. We outline such an outcome-based solution in Section 4.3. Specifically, in the credit setting, the outcome-based solution corresponds to giving out more loans to the protected group in a way that reduces profit for the bank compared to unconstrained profit maximization, but avoids loaning to those who are unlikely to benefit, resulting in a maximally improved group average credit score. The extent to which such a solution could form the basis of successful regulation depends on the accuracy of the available outcome model.

This raises the question if our model of outcomes is rich enough to faithfully capture realistic phenomena. By focusing on the impact that selection has on individuals at a given score, we model the effects for those *not* selected as zero-mean. For example, not getting a loan in our model has no negative effect on the credit score of an individual.¹ This does not mean that wrongful rejection (i.e., a false negative) has no visible manifestation in our model. If a classifier has a higher false negative rate in one group than in another, we expect the classifier to increase the disparity between the two groups (under natural assumptions). In other words, in our outcome-based model, the harm of denied opportunity manifests as growing disparity between the groups. The cost of a false negative could also be incorporated directly into the outcome-based model by a simple modification (see Footnote 2). This may be fitting in some applications where the immediate impact of a false negative to the individual is not zero-mean, but significantly reduces their future success probability.

In essence, the formalism we propose requires us to understand the two-variable causal mechanism that translates decisions to outcomes. This can be seen as relaxing the requirements compared with recent work on avoiding discrimination through causal reasoning that often required stronger assumptions [Kusner et al., 2017, Nabi and Shpitser, 2017, Kilbertus et al., 2017]. In particular, these works required knowledge of how sensitive attributes (such as gender, race, or proxies thereof)

¹In reality, a denied credit inquiry may lower one’s credit score, but the effect is small compared to a default event.

causally relate to various other variables in the data. Our model avoids the delicate modeling step involving the sensitive attribute, and instead focuses on an arguably more tangible economic mechanism. Nonetheless, depending on the application, such an understanding might necessitate greater domain knowledge and additional research into the specifics of the application. This is consistent with much scholarship that points to the context-sensitive nature of fairness in machine learning.

2 Problem Setting

We consider two *groups* A and B, which comprise a g_A and $g_B = 1 - g_A$ fraction of the total population, and an *institution* which makes a binary decision for each individual in each group, called *selection*. Individuals in each group are assigned *scores* in $\mathcal{X} := [C]$, and the scores for group $j \in \{A, B\}$ are distributed according $\pi_j \in \text{Simplex}^{C-1}$. The institution selects a *policy* $\tau := (\tau_A, \tau_B) \in [0, 1]^{2C}$, where $\tau_j(x)$ corresponds to the probability the institution selects an individual in group j with score x . One should think of a score as an abstract quantity which summarizes how well an individual is suited to being selected; examples are provided at the end of this section.

We assume that the institution is utility-maximizing, but may impose certain constraints to ensure that the policy τ is *fair*, in a sense described in Section 2.2. We assume that there exists a function $u : C \rightarrow \mathbb{R}$, such that the institution’s expected utility for a policy τ is given by

$$\mathcal{U}(\tau) = \sum_{j \in \{A, B\}} g_j \sum_{x \in \mathcal{X}} \tau_j(x) \pi_j(x) u(x). \quad (1)$$

Novel to this work, we focus on the effect of the selection policy τ on the groups A and B. We quantify these *outcomes* in terms of an average effect that a policy τ_j has on group j . Formally, for a function $\Delta(x) : \mathcal{X} \rightarrow \mathbb{R}$, we define the average change of the mean score μ_j for group j

$$\Delta\mu_j(\tau) := \sum_{x \in \mathcal{X}} \pi_j(x) \tau_j(x) \Delta(x). \quad (2)$$

We remark that many of our results also go through if $\Delta\mu_j(\tau)$ simply refers to an abstract change in well-being, not necessarily a change in the mean score. Furthermore, it is possible to modify the definition of $\Delta\mu_j(\tau)$ such that it directly considers outcomes of those who are not selected.² Lastly, we assume that the *success* of an individual is independent of their group given the score; that is, the score summarizes all relevant information about the success event, so there exists a function $\rho : \mathcal{X} \rightarrow [0, 1]$ such that individuals of score x succeed with probability $\rho(x)$.

We now introduce the specific domain of credit scores as a running example in the rest of the paper, after which we present two more examples showing the general applicability of our formulation to many domains.

Example 2.1 (Credit scores). In the setting of loans, scores $x \in [C]$ represent credit scores, and the bank serves as the institution. The bank chooses to grant or refuse loans to individuals according to a policy τ . Both bank and personal utilities are given as functions of loan repayment, and

² If we consider functions $\Delta_p(x) : \mathcal{X} \rightarrow \mathbb{R}$ and $\Delta_n(x) : \mathcal{X} \rightarrow \mathbb{R}$ to represent the average effect of selection and non-selection respectively, then $\Delta\mu_j(\tau) := \sum_{x \in \mathcal{X}} \pi_j(x) (\tau_j(x) \Delta_p(x) + (1 - \tau_j(x)) \Delta_n(x))$. This model corresponds to replacing $\Delta(x)$ in the original outcome definition with $\Delta_p(x) - \Delta_n(x)$, and adding a offset $\sum_{x \in \mathcal{X}} \pi_j(x) \Delta_n(x)$. Under the assumption that $\Delta_p(x) - \Delta_n(x)$ increases in x , this model gives rise to outcomes curves resembling those in Figure 1 up to vertical translation. All presented results hold unchanged under the further assumption that $\Delta\mu(\beta^{\text{MaxUtil}}) \geq 0$.

therefore depend on the success probabilities $\rho(x)$, representing the probability that any individual with credit score x can repay a loan within a fixed time frame. The expected utility to the bank is given by the expected return from a loan, which can be modeled as an affine function of $\rho(x)$: $u(x) = u_+\rho(x) + u_-(1 - \rho(x))$, where u_+ denotes the profit when loans are repaid and u_- the loss when they are defaulted on. Individual outcomes of being granted a loan are based on whether or not an individual repays the loan, and a simple model for $\Delta(x)$ may also be affine in $\rho(x)$: $\Delta(x) = c_+\rho(x) + c_-(1 - \rho(x))$, modified accordingly at boundary states. The constant c_+ denotes the gain in credit score if loans are repaid and c_- is the score penalty in case of default.

Example 2.2 (Advertising). A second illustrative example is given by the case of advertising agencies making decisions about which groups to target. An individual with product interest score x responds positively to an ad with probability $\rho(x)$. The ad agency experiences utility $u(x)$ related to click-through rates, which increases with $\rho(x)$. Individuals who see the ad but are uninterested may react negatively (becoming less interested in the product), and $\Delta(x)$ encodes the interest change. If the product is a positive good like education or employment opportunities, interest can correspond to well-being. Thus the advertising agency’s incentives to only show ads to individuals with extremely high interest may leave behind groups whose interest is lower on average. A related historical example occurred in advertisements for computers in the 1980s, where male consumers were targeted over female consumers, arguably contributing to the current gender gap in computing.

Example 2.3 (College Admissions). The scenario of college admissions or scholarship allotments can also be considered within our framework. Colleges may select certain applicants for acceptance according to a score x , which could be thought encode a “college preparedness” measure. The students who are admitted might “succeed” (this could be interpreted as graduating, graduating with honors, finding a job placement, etc.) with some probability $\rho(x)$ depending on their preparedness. The college might experience a utility $u(x)$ corresponding to alumni donations, or positive rating when a student succeeds; they might also show a drop in rating or a loss of invested scholarship money when a student is unsuccessful. The student’s success in college will affect their later success, which could be modeled generally by $\Delta(x)$. In this scenario, it is challenging to ensure that a single summary statistic x captures enough information about a student; it may be more appropriate to consider x as a vector as well as more complex forms of $\rho(x)$.

While a variety of applications are modeled faithfully within our framework, there are limitations to the accuracy with which real-life phenomenon can be measured by strictly binary decisions and success probabilities. Such binary rules are necessary for the definition and execution of existing fairness criteria, (see Sec. 2.2) and as we will see, even modeling these facets of decision making as binary allows for complex and interesting behavior.

2.1 The Outcome Curve

We now introduce important outcome regimes, stated in terms of the change in average group score. A policy (τ_A, τ_B) is said to cause *active harm* to group j if $\Delta\mu_j(\tau_j) < 0$, *stagnation* if $\Delta\mu_j(\tau_j) = 0$, and *improvement* if $\Delta\mu_j(\tau_j) > 0$. Under our model, **MaxUtil** policies can be chosen in a standard fashion which applies the same threshold τ^{MaxUtil} for both groups, and is agnostic to the distributions π_A and π_B . Hence, if we define

$$\Delta\mu_j^{\text{MaxUtil}} := \Delta\mu_j(\tau^{\text{MaxUtil}}) \quad (3)$$




Figure 1: The above figure shows the *outcome curve*. The horizontal axis represents the selection rate for the population; the vertical axis represents the mean change in score. (a) depicts the full spectrum of outcome regimes, and colors indicate regions of active harm, relative harm, and no harm. In (b): a group that has much potential for gain, in (c): a group that has no potential for gain.

we say that a policy causes *relative harm* to group j if $\Delta\mu_j(\tau_j) < \Delta\mu_j^{\text{MaxUtil}}$, and *relative improvement* if $\Delta\mu_j(\tau_j) > \Delta\mu_j^{\text{MaxUtil}}$. In particular, we focus on these outcomes for a disadvantaged group, and consider whether imposing a fairness constraint improves their outcomes relative to the `MaxUtil` strategy. From this point forward, we take A to be disadvantaged or protected group.

Figure 1 displays the important outcome regimes in terms of *selection rates* $\beta_j := \sum_{x \in \mathcal{X}} \pi_j(x) \tau_j(x)$. This succinct characterization is possible when considering decision rules based on (possibly randomized) score thresholding, in which all individuals with scores above a threshold are selected. In Section 5, we justify the restriction to such *threshold policies* by showing it preserves optimality. In Section 5.1, we show that the outcome curve is concave, thus implying that it takes the shape depicted in Figure 1. To explicitly connect selection rates to decision policies, we define the rate function $r_\pi(\tau_j)$ which returns the proportion of group j selected by the policy. We show that this function is invertible for a suitable class of threshold policies, and in fact the outcome curve is precisely the graph of the map from selection rate to outcome $\beta \mapsto \Delta\mu_A(r_{\pi_A}^{-1}(\beta))$. Next, we define the values of β that mark boundaries of the outcome regions.

Definition 2.1 (Selection rates of interest). Given the protected group A , the following selection rates are of interest in distinguishing between qualitatively different classes of outcomes (Figure 1). We define β^{MaxUtil} as the selection rate for A under `MaxUtil`; β_0 as the harm threshold, such that $\Delta\mu_A(r_{\pi_A}^{-1}(\beta_0)) = 0$; β^* as the selection rate such that $\Delta\mu_A$ is maximized; $\bar{\beta}$ as the outcome-complement of the `MaxUtil` selection rate, $\Delta\mu_A(r_{\pi_A}^{-1}(\bar{\beta})) = \Delta\mu_A(r_{\pi_A}^{-1}(\beta^{\text{MaxUtil}}))$ with $\bar{\beta} > \beta^{\text{MaxUtil}}$.

2.2 Decision Rules and Fairness Criteria

We will consider policies that maximize the institution’s total expected utility, potentially subject to a constraint: $\tau \in \mathcal{C} \subset [0, 1]^{2C}$ which enforces some notion of “fairness”. Formally, the institution selects $\tau_* \in \operatorname{argmax} \mathcal{U}(\tau)$ s.t. $\tau \in \mathcal{C}$. We consider the three following constraints:

Definition 2.2 (Fairness criteria). The *maximum utility* (**MaxUtil**) policy corresponds to the null-constraint $\mathcal{C} = [0, 1]^{2C}$, so that the institution is free to focus solely on utility. The *demographic parity* (**DemParity**) policy results in equal selection rates between both groups. Formally, the constraint is $\mathcal{C} = \{(\tau_A, \tau_B) : \sum_{x \in \mathcal{X}} \pi_A(x) \tau_A = \sum_{x \in \mathcal{X}} \pi_B(x) \tau_B\}$. The *equal opportunity* (**EqOpt**) policy results in equal true positive rates (TPR) between both group, where TPR is defined as $\operatorname{TPR}_j(\tau) := \frac{\sum_{x \in \mathcal{X}} \pi_j(x) \rho(x) \tau(x)}{\sum_{x \in \mathcal{X}} \pi_j(x) \rho(x)}$. **EqOpt** ensures that the conditional probability of selection given that the individual will be successful is independent of the population, formally enforced by the constraint $\mathcal{C} = \{(\tau_A, \tau_B) : \operatorname{TPR}_A(\tau_A) = \operatorname{TPR}_B(\tau_B)\}$.

Just as the expected outcome $\Delta\mu$ can be expressed in terms of selection rate for threshold policies, so can the total utility $\mathcal{U}$. In the unconstrained case, $\mathcal{U}$ varies independently over the selection rates for group A and B; however, in the presence of fairness constraints the selection rate for one group determines the allowable selection rate for the other. The selection rates must be equal for **DemParity**, but for **EqOpt** we can define a *transfer function*, $G^{(A \rightarrow B)}$, which for every loan rate β in group A gives the loan rate in group B that has the same true positive rate. Therefore, when considering threshold policies, decision rules amount to maximizing functions of single parameters. This idea is expressed in Figure 2, and underpins the results to follow.

3 Results

In order to clearly characterize the outcome of applying fairness constraints, we make the following assumption.

Assumption 1 (Institution utilities). *The institution’s individual utility function is more stringent than the expected score changes, $u(x) > 0 \implies \Delta(x) > 0$. (For the linear form presented in Example 2.1, $\frac{u_-}{u_+} < \frac{c_-}{c_+}$ is necessary and sufficient.)*

This simplifying assumption quantifies the intuitive notion that institutions take a greater risk by accepting than the individual does by applying. For example, in the credit setting, a bank loses the amount loaned in the case of a default, but makes only interest in case of a payback. Using Assumption 1, we can restrict the position of **MaxUtil** on the outcome curve in the following sense.

Proposition 3.1 (**MaxUtil** does not cause active harm). *Under Assumption 1, $0 \leq \Delta\mu^{\operatorname{MaxUtil}} \leq \Delta\mu^*$.*

We direct the reader to Appendix C for the proof of the above proposition, and all subsequent results presented in this section. The results are corollaries to theorems presented in Section 6.

3.1 Prospects and Pitfalls of Fairness Criteria

We begin by characterizing general settings under which fairness criteria act to improve outcomes over unconstrained **MaxUtil** strategies. For this result, we will assume that group A is disadvantaged




Figure 2: Both outcomes $\Delta\mu$ and institution utilities $\mathcal{U}$ can be plotted as a function of selection rate for one group. The maxima of the utility curves determine the selection rates resulting from various decision rules.

in the sense that the **MaxUtil** acceptance rate for B is large compared to relevant acceptance rates for A.

Corollary 3.2 (Fairness Criteria can cause Relative Improvement). *(a) Under the assumption that $\beta_A^{\text{MaxUtil}} < \bar{\beta}$ and $\beta_B^{\text{MaxUtil}} > \beta_A^{\text{MaxUtil}}$, there exist population proportions $g_0 < g_1 < 1$ such that, for all $g_A \in [g_0, g_1]$, $\beta_A^{\text{MaxUtil}} < \beta_A^{\text{DemParity}} < \bar{\beta}$. That is, **DemParity** causes relative improvement.*

*(b) Under the assumption that there exist $\beta_A^{\text{MaxUtil}} < \beta < \beta' < \bar{\beta}$ such that $\beta_B^{\text{MaxUtil}} > G^{(A \rightarrow B)}(\beta)$, $G^{(A \rightarrow B)}(\beta')$, there exist population proportions $g_2 < g_3 < 1$ such that, for all $g_A \in [g_2, g_3]$, $\beta_A^{\text{MaxUtil}} < \beta_A^{\text{EqOpt}} < \bar{\beta}$. That is, **EqOpt** causes relative improvement.*

This result gives the conditions under which we can guarantee the existence of settings in which fairness criteria cause improvement relative to **MaxUtil**. Relying on machinery proved in Section 6, the result follows from comparing the position of optima on the utility curve to the outcome curve. Figure 2 displays an illustrative example of both the outcome curve and the institutions' utility $\mathcal{U}$ as a function of the selection rates in group A. In the utility function (1), the contributions of each group are weighted by their population proportions g_j , and thus the resulting selection rates are sensitive to these proportions.

As we see in the remainder of this section, fairness criteria can achieve nearly any position along the outcome curve under the right conditions. This fact comes from the potential mismatch between the outcomes, controlled by Δ , and the institution's utility u .

The next theorem implies that **DemParity** can be bad for long term well-being of the protected group by being over-generous, under the mild assumption that $\Delta\mu_A(\beta_B^{\text{MaxUtil}}) < 0$:

Corollary 3.3 (**DemParity** can cause harm by being over-eager). *Fix a selection rate β . Assume that $\beta_B^{\text{MaxUtil}} > \beta > \beta_A^{\text{MaxUtil}}$. Then, there exists a population proportion g_0 such that, for all $g_A \in [0, g_0]$, $\beta_A^{\text{DemParity}} > \beta$. In particular, when $\beta = \beta_0$, **DemParity** causes active harm, and when $\beta = \bar{\beta}$, **DemParity** causes relative harm.*

The assumption $\Delta\mu_A(\beta_B^{\text{MaxUtil}}) < 0$ implies that a policy which selects individuals from group A at the selection rate that **MaxUtil** would have used for group B necessarily lowers average score in A. This is one natural notion of protected group A's ‘disadvantage’ relative to group B. In this case, **DemParity** penalizes the scores of group A even more than a naive **MaxUtil** policy, as long as group proportion g_A is small enough. Again, small g_A is another notion of group disadvantage.

Using credit scores as an example, Corollary 3.3 tells us that an overly aggressive fairness criterion will give too many loans to people in a protected group who cannot pay them back, hurting the group’s credit scores on average. In the following theorem, we show that an analogous result holds for **EqOpt**.

Corollary 3.4 (**EqOpt** can cause harm by being over-eager). *Suppose that $\beta_B^{\text{MaxUtil}} > G^{(A \rightarrow B)}(\beta)$ and $\beta > \beta_A^{\text{MaxUtil}}$. Then, there exists a population proportion g_0 such that, for all $g_A \in [0, g_0]$, $\beta_A^{\text{EqOpt}} > \beta$. In particular, when $\beta = \beta_0$, **EqOpt** causes active harm, and when $\beta = \bar{\beta}$, **EqOpt** causes relative harm.*

We remark that in Corollary 3.4, we rely on the *transfer function*, $G^{(A \rightarrow B)}$, which for every loan rate β in group A gives the loan rate in group B that has the same true positive rate. Notice that if $G^{(A \rightarrow B)}$ were the identity function, Corollary 3.3 and Corollary 3.4 would be exactly the same. Indeed, our framework (detailed in Section 6 and Appendix B) unifies the analyses for a large class of fairness constraints that includes **DemParity** and **EqOpt** as specific cases, and allows us to derive results about impact on $\Delta\mu$ using general techniques. In the next section, we present further results that compare the fairness criteria, demonstrating the usefulness of our technical framework.

3.2 Comparing EqOpt and DemParity

Our analysis of the acceptance rates of **EqOpt** and **DemParity** in Section 6 suggests that it is difficult to compare **DemParity** and **EqOpt** without knowing the full distributions π_A, π_B , which is necessary to compute the transfer function $G^{(A \rightarrow B)}$. In fact, we have found that settings exist both in which **DemParity** causes harm while **EqOpt** causes improvement and in which **DemParity** causes improvement while **EqOpt** causes harm. There cannot be one general rule as to which fairness criteria provides better outcomes in all settings. We now present simple sufficient conditions on the geometry of the distributions for which **EqOpt** is always better than **DemParity** in terms of $\Delta\mu_A$.

Corollary 3.5 (**EqOpt** may avoid active harm where **DemParity** fails). *Fix a selection rate β . Suppose π_A, π_B are identical up to a translation with $\mu_A < \mu_B$, i.e. $\pi_A(x) = \pi_B(x + (\mu_B - \mu_A))$. For simplicity, take $\rho(x)$ to be linear in x . Suppose*

$$\beta > \sum_{x > \mu_A} \pi_A.$$

*Then there exists an interval $[g_1, g_2] \subseteq [0, 1]$, such that $\forall g_A > g_1, \beta^{\text{EqOpt}} < \beta$ while $\forall g_A < g_2, \beta^{\text{DemParity}} > \beta$. In particular, when $\beta = \beta_0$, this implies **DemParity** causes active harm but **EqOpt** causes improvement for $g_A \in [g_1, g_2]$, but for any g_A such that **DemParity** causes improvement, **EqOpt** also causes improvement.*

To interpret the conditions under which Corollary 3.5 holds, consider when we might have $\beta_0 > \sum_{x > \mu_A} \pi_A$. This is precisely when $\Delta\mu_A(\sum_{x > \mu_A} \pi_A) > 0$, that is, $\Delta\mu_A > 0$ for a policy that selects every individual whose score is above the group A mean, which is reasonable in reality.

Indeed, the converse would imply that group A has such low scores that even selecting all above average individuals in A would hurt the average score. In such a case, Corollary 3.5 suggests that EqOpt is better than DemParity at avoiding active harm, because it is more conservative. A natural question then is: can EqOpt cause relative harm by being too stingy?

Corollary 3.6 (DemParity never loans less than MaxUtil, but EqOpt might). *Recall the definition of the TPR functions TPR_j , and suppose that the MaxUtil policy τ^{MaxUtil} is such that*

$$\beta_A^{\text{MaxUtil}} < \beta_B^{\text{MaxUtil}} \text{ and } \text{TPR}_A(\tau^{\text{MaxUtil}}) > \text{TPR}_B(\tau^{\text{MaxUtil}}) \quad (4)$$

Then $\beta_A^{\text{EqOpt}} < \beta_A^{\text{MaxUtil}} < \beta_A^{\text{DemParity}}$. That is, EqOpt causes relative harm by selecting at a rate lower than MaxUtil.

The above theorem shows that DemParity is never stingier than MaxUtil to the protected group A, as long as A is disadvantaged in the sense that MaxUtil selects a larger proportion of B than A. On the other hand, EqOpt can select less of group A than MaxUtil, and by definition, cause relative harm. This is a surprising result about EqOpt, and this phenomenon arises from high levels of in-group inequality for group A. Moreover, we show in Appendix C that there are parameter settings where the conditions in Corollary 3.6 are satisfied even under a stringent notion of disadvantage we call CDF domination, described therein.

4 Relaxations of Constrained Fairness

4.1 Regularized fairness

In many cases, it may be unrealistic for an institution to ensure that fairness constraints are met exactly. However, one can consider “soft” formulations of fairness constraints which either penalized the differences in acceptance rate (DemParity) or the differences in TPR (EqOpt). In Appendix B, we formulate these soft constraints as regularized objectives. For example, a soft-DemParity can be rendered as

$$\max_{\tau := \tau_A, \tau_B} \mathcal{U}(\tau) - \lambda \Phi(\langle \pi_A, \tau_A \rangle - \langle \pi_B, \tau_B \rangle), \quad (5)$$

where $\lambda > 0$ is a regularization parameter, and $\Phi(t)$ is a convex regularization function. We show that the solutions to these objectives are threshold policies, and can be fully characterized in terms of the group-wise selection rate. We also make rigorous the notion that policies which solve the soft-constraint objective interpolate between MaxUtil policies at $\lambda = 0$ and hard-constrained policies (DemParity or EqOpt) as $\lambda \rightarrow \infty$. This fact is clearly demonstrated by the form of the solutions in the special case of the regularization function $\Phi(t) = |t|$, provided in the appendix.

4.2 Fairness Under Measurement Error

Next, consider the implications of an institution with imperfect knowledge of scores. Under a simple model in which the estimate of an individual’s score $X \sim \pi$ is prone to errors $e(X)$ such that $X + e(X) := \hat{X} \sim \hat{\pi}$. Constraining the error to be negative results in the setting that scores are systematically *underestimated*. In this setting, it is equivalent to consider the CDF of underestimated distribution $\hat{\pi}$ to be *dominated* by the CDF true distribution π , that is

$\sum_{x \geq c} \hat{\pi}(x) \leq \sum_{x \geq c} \pi(x)$ for all $c \in [C]$. Then we can compare the institution's behavior under this estimation to its behavior under the truth.

Proposition 4.1 (Underestimation causes underselection). *Fix the distribution of B as π_B and let β be the acceptance rate of A when the institution makes the decision using perfect knowledge of the distribution π_A . Denote $\hat{\beta}$ as the acceptance rate when the group is instead taken as $\hat{\pi}_A$. Then $\beta_A^{\text{MaxUtil}} > \hat{\beta}_A^{\text{MaxUtil}}$ and $\beta_A^{\text{DemParity}} > \hat{\beta}_A^{\text{DemParity}}$. If the errors are further such that the true TPR dominates the estimated TPR, it is also true that $\beta_A^{\text{EqOpt}} > \hat{\beta}_A^{\text{EqOpt}}$.*

Because fairness criteria encourage a higher selection rate for disadvantaged groups (Corollary 3.2), systematic underestimation widens the regime of their applicability. Furthermore, since the estimated **MaxUtil** policy underloans, the region for relative improvement in the outcome curve (Figure 1) is larger, corresponding to more regimes under which fairness criteria can yield favorable outcomes. Thus the potential for measurement error should be a factor when motivating these criteria.

4.3 Outcome-based alternative

As explained in the preceding sections, fairness criteria may actively harm disadvantaged groups. It is thus natural to consider a modified decision rule which involves the explicit maximization of $\Delta\mu_A$. In this case, imagine that the institution's primary goal is to aid the disadvantaged group, subject to a limited profit loss compared to the maximum possible expected profit $\mathcal{U}^{\text{MaxUtil}}$. The corresponding problem is as follows.

$$\max_{\tau_A} \Delta\mu_A(\tau_A) \quad \text{s.t.} \quad \mathcal{U}_A^{\text{MaxUtil}} - \mathcal{U}(\tau) < \delta. \quad (6)$$

Unlike the fairness constrained objective, this objective no longer depends on group B and instead depends on our model of the mean score change in group A , $\Delta\mu_A$.

Proposition 4.2 (Outcome-based solution). *In the above setting, the optimal bank policy τ_A is a threshold policy with selection rate $\beta = \min\{\beta^*, \beta^{\text{max}}\}$, where β^* is the outcome-optimal loan rate and β^{max} is the maximum loan rate under the bank's "budget".*

The above formulation's advantage over fairness constraints is that it directly optimizes the outcome of A and can be approximately implemented given reasonable ability to predict outcomes. Importantly, this objective shifts the focus to outcome modeling, highlighting the importance of domain specific knowledge. Future work can consider strategies that are robust to outcome model errors.

5 Optimality of Threshold Policies

Next, we move towards statements of the main theorems underlying the results presented in Section 3. We begin by establishing notation which we shall use throughout. Recall that $\circ$ denotes the Hadamard product between vectors. We identify functions mapping $\mathcal{X} \rightarrow \mathbb{R}$ with vectors in $\mathbb{R}^C$. We also define the group-wise utilities

$$\mathcal{U}_j(\tau_j) := \sum_{x \in \mathcal{X}} \pi_j(x) \tau_j(x) u(x), \quad (7)$$

so that for $\tau = (\tau_A, \tau_B)$, $\mathcal{U}(\tau) := g_A \mathcal{U}_A(\tau_A) + g_B \mathcal{U}_B(\tau_B)$.

First, we formally describe threshold policies, and rigorously justify why we may always assume without loss of generality that the institution adopts policies of this form.

Definition 5.1 (Threshold selection policy). A single group selection policy $\tau \in [0, 1]^C$ is called a *threshold policy* if it has the form of a randomized threshold on score:

$$\tau_{c,\gamma} = \begin{cases} 1, & x > c \\ \gamma, & x = c \\ 0, & x < c \end{cases}, \text{ for some } c \in [C] \text{ and } \gamma \in (0, 1]. \quad (8)$$

As a technicality, if no members of a population have a given score $x \in \mathcal{X}$, there may be multiple threshold policies which yield equivalent selection rates for a given population. To avoid redundancy, we introduce the notation $\tau_j \cong_{\pi_j} \tau'_j$ to mean that the set of scores on which τ_j and τ'_j differ has probability 0 under π_j ; formally, $\sum_{x: \tau_j(x) \neq \tau'_j(x)} \pi_j(x) = 0$. For any distribution π_j , $\cong_{\pi_j}$ is an equivalence relation. Moreover, we see that if $\tau_j \cong_{\pi_j} \tau'_j$, then τ_j and τ'_j both provide the same utility for the institution, induce the same outcomes for individuals in group j , and have the same selection and true positive rates. Hence, if (τ_A, τ_B) is an optimal solution to any of **MaxUtil**, **EqOpt**, or **DemParity**, so is any (τ'_A, τ'_B) for which $\tau_A \cong_{\pi_A} \tau'_A$ and $\tau_B \cong_{\pi_B} \tau'_B$.

For threshold policies in particular, their equivalence class under $\cong_{\pi_j}$ is uniquely determined by the selection rate function,

$$r_{\pi_j}(\tau_j) := \sum_{x \in \mathcal{X}} \pi_j(x) \tau_j(x), \quad (9)$$

which denotes the fraction of group j which is selected. Indeed, we have the following lemma (proved in Appendix A.1):

Lemma 5.1. *Let τ_j and τ'_j be threshold policies. Then $\tau_j \cong_{\pi_j} \tau'_j$ if and only if $r_{\pi_j}(\tau_j) = r_{\pi_j}(\tau'_j)$. Further, $r_{\pi_j}(\tau_j)$ is a bijection from $\mathcal{T}_{\text{thresh}}(\pi_j)$ to $[0, 1]$, where $\mathcal{T}_{\text{thresh}}(\pi_j)$ is the set of equivalence classes between threshold policies under $\cong_{\pi_j}$. Finally, $\pi_j \circ r_{\pi_j}^{-1}(\beta_j)$ is well defined.*

Remark that $r_{\pi_j}^{-1}(\beta_j)$ is an equivalence class rather than a single policy. However, $\pi_j \circ r_{\pi_j}^{-1}(\tau_j)$ is well defined, meaning that $\pi_j \circ \tau_j = \pi_j \circ \tau'_j$ for any two policies in the same equivalence class. Since all quantities of interest will only depend on policies τ_j through $\pi_j \circ \tau_j$, it does not matter *which* representative of $r_{\pi_j}^{-1}(\beta_j)$ we pick. Hence, abusing notation slightly, we shall represent $\mathcal{T}_{\text{thresh}}(\pi_j)$ by choosing one representative from each equivalence class under $\cong_{\pi_j}$ ³.

It turns out the policies which arise in this way are always optimal in the sense that, for a given loan rate β_j , the threshold policy $r_{\pi_j}^{-1}(\beta_j)$ is the (essentially unique) policy which maximizes both the institution's utility and the utility of the group. Defining the group-wise utility,

$$\mathcal{U}_j(\tau_j) := \sum_{x \in \mathcal{X}} \mathbf{u}(x) \pi_j(x) \tau_j(x), \quad (10)$$

we have the following result:

³One way to do this is to consider the set of all threshold policies $\tau_{c,\gamma}$ such that, $\gamma = 1$ if $\pi_j(c) = 0$ and $\pi_j(c-1) > 0$ if $\gamma = 1$ and $c > 1$.

Proposition 5.1 (Threshold policies are preferable). *Suppose that $\mathbf{u}(x)$ and $\Delta(x)$ are strictly increasing in x . Given any loaning policy τ_j for population with distribution π_j , then the policy $\tau_j^{\text{thresh}} := r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j)) \in \mathcal{T}_{\text{thresh}}(\pi_j)$ satisfies*

$$\Delta\mu_j(\tau_j^{\text{thresh}}) \geq \Delta\mu_j(\tau_j) \text{ and } \mathcal{U}_j(\tau_j^{\text{thresh}}) \geq \mathcal{U}_j(\tau_j). \quad (11)$$

Moreover, both inequalities hold with equality if and only if $\tau_j \cong_{\pi_j} \tau_j^{\text{thresh}}$.

The map $\tau_j \mapsto r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$ can be thought of transforming an arbitrary policy τ_j into a threshold policy with the same selection rate. In this language, the above proposition states that this map never reduces institution utility or individual outcomes. We can also show that optimal **MaxUtil** and **DemParity** policies are threshold policies, as well as all **EqOpt** policies under an additional assumption:

Proposition 5.2 (Existence of optimal threshold policies under fairness constraints). *Suppose that $\mathbf{u}(x)$ is strictly increasing in x . Then all optimal **MaxUtil** policies (τ_A, τ_B) satisfy $\tau_j \cong_{\pi_j} r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$ for $j \in \{A, B\}$. The same holds for all optimal **DemParity** policies, and if in addition $\mathbf{u}(x)/\rho(x)$ is increasing, the same is true for all optimal **EqOpt** policies.*

To prove proposition 5.1, we invoke the following general lemma which is proved using standard convex analysis arguments (in Appendix A.2):

Lemma 5.2. *Let $\mathbf{v} \in \mathbb{R}^C$, and let $\mathbf{w} \in \mathbb{R}_{>0}^C$, and suppose either that $\mathbf{v}(x)$ is increasing in x , and $\mathbf{v}(x)/\mathbf{w}(x)$ is increasing or, $\forall x \in \mathcal{X}$, $\mathbf{w}(x) = 0$. Let $\pi \in \text{Simplex}^{C-1}$ and fix $t \in [0, \sum_{x \in \mathcal{X}} \pi(x) \cdot \mathbf{w}(x)]$. Then any*

$$\tau^* \in \arg \max_{\tau \in [0,1]^C} \langle \mathbf{v} \circ \pi, \tau \rangle \quad \text{s.t.} \quad \langle \pi \circ \mathbf{w}, \tau \rangle = t \quad (12)$$

satisfies $\tau^* \cong_{\pi} r_{\pi}^{-1}(r_{\pi}(\tau^*))$. Moreover, at least one maximizer $\tau^* \in \mathcal{T}_{\text{thresh}}(\pi)$ exists.

Proof of Proposition 5.1. We will first prove Proposition 5.1 for the function $\mathcal{U}_j$. Given our nominal policy τ_j , let $\beta_j = r_{\pi_j}(\tau_j)$. We now apply Lemma 5.2 with $\mathbf{v}(x) = \mathbf{u}(x)$ and $\mathbf{w}(x) = 1$. For this choice of $\mathbf{v}$ and $\mathbf{w}$, $\langle \mathbf{v}, \tau \rangle = \mathcal{U}_j(\tau)$ and that $\langle \pi_j \circ \mathbf{w}, \tau \rangle = r_{\pi_j}(\tau)$. Then, if $\tau_j \in \arg \max_{\tau} \mathcal{U}_j(\tau)$ s.t. $r_{\pi_j}(\tau) = \beta_j$, Lemma 12 implies that $\tau_j \cong_{\pi_j} r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$.

On the other hand, assume that $\tau_j \cong_{\pi_j} r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$. We show that $r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$ is a maximizer; which will imply that τ_j is a maximizer since $\tau_j \cong_{\pi_j} r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$ implies that $\mathcal{U}_j(\tau_j) = \mathcal{U}_j(r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j)))$. By Lemma 5.2 there exists a maximizer $\tau_j^* \in \mathcal{T}_{\text{thresh}}(\pi)$, which means that $\tau_j^* = r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j^*))$. Since τ_j^* is feasible, we must have $r_{\pi_j}(\tau_j^*) = r_{\pi_j}(\tau_j)$, and thus $\tau_j^* = r_{\pi_j}^{-1}(r_{\pi_j}(\tau_j))$, as needed. The same argument follows verbatim if we instead choose $\mathbf{v}(x) = \Delta(x)$, and compute $\langle \mathbf{v}, \tau \rangle = \Delta\mu_j(\tau)$. $\square$

We now argue Proposition 5.2 for **MaxUtil**, as it is a straightforward application of Lemma 5.2. We will prove Proposition 5.2 for **DemParity** and **EqOpt** separately in Sections 6.1 and 6.2.

Proof of Proposition 5.2 for MaxUtil. **MaxUtil** follows from lemma 5.2 with $\mathbf{v}(x) = \mathbf{u}(x)$, and $t = 0$ and $\mathbf{w} = \mathbf{0}$. $\square$

5.1 Quantiles and Concavity of the Outcome Curve

To further our analysis, we now introduce left and right quantile functions, allowing us to specify thresholds in terms of both selection rate and score cutoffs.

Definition 5.2 (Upper quantile function). Define Q to be the upper quantile function corresponding to π , i.e.

$$Q_j(\beta) = \operatorname{argmax}\{c : \sum_{x=c}^C \pi_j(x) > \beta\} \quad \text{and} \quad Q_j^+(\beta) := \operatorname{argmax}\{c : \sum_{x=c}^C \pi_j(x) \geq \beta\}. \quad (13)$$

Crucially $Q(\beta)$ is continuous from the right, and $Q^+(\beta)$ is continuous from the left. Further, $Q(\cdot)$ and $Q^+(\cdot)$ allow us to compute derivatives of key functions, like the mapping from selection rate β to the group outcome associated with a policy of that rate, $\Delta\mu(r_\pi^{-1}(\beta))$. Because we take π to have discrete support, all functions in this work are *piecewise linear*, so we shall need to distinguish between the left and right derivatives, defined as follows

$$\partial_- f(x) := \lim_{t \rightarrow 0^-} \frac{f(x+t) - f(x)}{t} \quad \text{and} \quad \partial_+ f(y) := \lim_{t \rightarrow 0^+} \frac{f(y+t) - f(y)}{t}. \quad (14)$$

For f supported on $[a, b]$, we say that f is left- (resp. right-) differentiable if $\partial_- f(x)$ exists for all $x \in (a, b]$ (resp. $\partial_+ f(y)$ exists for all $y \in [a, b)$). We now state the fundamental derivative computation which underpins the results to follow:

Lemma 5.3. *Let e_x denote the vector such that $e_x(x) = 1$, and $e_x(x') = 0$ for $x' \neq x$. Then $\pi_j \circ r_{\pi_j}^{-1}(\beta) : [0, 1] \rightarrow [0, 1]^C$ is continuous, and has left and right derivatives*

$$\partial_+ \left(\pi_j \circ r_{\pi_j}^{-1}(\beta) \right) = e_{Q(\beta)} \quad \text{and} \quad \partial_- \left(\pi_j \circ r_{\pi_j}^{-1}(\beta) \right) = e_{Q^+(\beta)}. \quad (15)$$

The above lemma is proved in Appendix A.3. Moreover, Lemma 5.3 implies that the outcome curve is concave under the assumption that $\Delta(x)$ is monotone:

Proposition 5.3. *Let π be a distribution over C states. Then $\beta \mapsto \Delta\mu(r_\pi^{-1}(\beta))$ is concave. In fact, if $w(x)$ is any non-decreasing map from $\mathcal{X} \rightarrow \mathbb{R}$, $\beta \mapsto \langle w, r_\pi^{-1}(\beta) \rangle$ is concave.*

Proof. Recall that a univariate function f is concave (and finite) on $[a, b]$ if and only (a) f is left- and right-differentiable, (b) for all $x \in (a, b)$, $\partial_- f(x) \geq \partial_+ f(x)$ and (c) for any $x > y$, $\partial_- f(x) \leq \partial_+ f(y)$.

Observe that $\Delta\mu(r_\pi^{-1}(\beta)) = \langle \Delta, \pi \circ r_\pi^{-1}(\beta) \rangle$. By Lemma 5.3, $\pi \circ r_\pi^{-1}(\beta)$ has right and left derivatives $e_{Q(\beta)}$ and $e_{Q^+(\beta)}$. Hence, we have that

$$\partial_+ \Delta\mu(\beta_B) = \Delta(Q(\beta_B)) \quad \text{and} \quad \partial_- \Delta\mu(\beta_B) = \Delta(Q^+(\beta_B)). \quad (16)$$

Using the fact that $\Delta(x)$ is monotone, and that $Q \leq Q^+$, we see that $\partial_+ \Delta\mu(f_\pi^{-1}(\beta_B)) \leq \partial_- \Delta\mu(f_\pi^{-1}(\beta_B))$, and that $\partial_- \Delta\mu(f_\pi^{-1}(\beta_B))$ and $\partial_+ \Delta\mu(f_\pi^{-1}(\beta_B))$ are non-increasing, from which it follows that $\Delta\mu(f_\pi^{-1}(\beta_B))$ is concave. The general concavity result holds by replacing $\Delta(x)$ with $w(x)$. $\square$




Figure 3: Considering the utility as a function of selection rates, fairness constraints correspond to restricting the optimization to one-dimensional curves. The **DemParity** (DP) constraint is a straight line with slope 1, while the **EqOpt** (EO) constraint is a curve given by the graph of $G^{(A \rightarrow B)}$. The derivatives considered throughout Section 6 are taken with respect to the selection rate β_A (horizontal axis); projecting the EO and DP constraint curves to the horizontal axis recovers concave utility curves such as those shown in the lower panel of Figure 2 (where **MaxUtil** in is represented by a horizontal line through the MU optimal solution).

6 Proofs of Main Theorems

We are now ready to present and prove theorems that characterize the selection rates under fairness constraints, namely **DemParity** and **EqOpt**. These characterizations are crucial for proving the results in Section 3. Our computations also generalize readily to other linear constraints, in a way that will become clear in Section 6.2.

6.1 A Characterization Theorem for DemParity

In this section, we provide a theorem that gives an explicit characterization for the range of selection rates β_A for **A** when the bank loans according to **DemParity**. Observe that the **DemParity** objective corresponds to solving the following linear program:

$$\max_{\tau=(\tau_A, \tau_B) \in [0,1]^{2C}} \mathcal{U}(\tau) \quad \text{s.t.} \quad \langle \pi_A, \tau_A \rangle = \langle \pi_B, \tau_B \rangle.$$

Let us introduce the auxiliary variable $\beta := \langle \pi_A, \tau_A \rangle = \langle \pi_B, \tau_B \rangle$ corresponding to the selection rate which is held constant across groups, so that all feasible solutions lie on the green DP line in Figure 3. We can then express the following equivalent linear program:

$$\max_{\tau=(\tau_A, \tau_B) \in [0,1]^{2C}, \beta \in [0,1]} \mathcal{U}(\tau) \quad \text{s.t.} \quad \beta = \langle \pi_j, \tau_j \rangle, \quad j \in \{A, B\}.$$

This is equivalent because, for a given β , Proposition 5.2 says that the utility maximizing policies are of the form $\tau_j = r_{\pi_j}^{-1}(\beta)$. We now prove this:

Proof of Proposition 5.2 for DemParity. Noting that $r_{\pi_j}(\tau_j) = \langle \pi_j, \tau_j \rangle$, we see that, by Lemma 5.2, under the special case where $\mathbf{v}(x) = \mathbf{u}(x)$ and $\mathbf{w}(x) = 1$, the optimal solution $(\tau_A^*(\beta), \tau_B^*(\beta))$ for fixed $r_{\pi_A}(\tau_A) = r_{\pi_B}(\tau_B) = \beta$ can be chosen to coincide with the threshold policies. Optimizing over β , the global optimal must coincide with thresholds. $\square$

Hence, any optimal policy is equivalent to the threshold policy $\tau = (r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta))$, where β solves the following optimization:

$$\max_{\beta \in [0,1]} \mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta))) . \quad (17)$$

We shall show that the above expression is in fact a *concave* function in β , and hence the set of optimal selection rates can be characterized by first order conditions. This is presented formally in the following theorem:

Theorem 6.1 (Selection rates for DemParity). *The set of optimal selection rates β^* satisfying (17) forms a continuous interval $[\beta_{\text{DemParity}}^-, \beta_{\text{DemParity}}^+]$, such that for any $\beta \in [0, 1]$, we have*

$$\begin{aligned} \beta < \beta_{\text{DemParity}}^- & \text{ if } g_A \mathbf{u}(Q_A(\beta)) + g_B \mathbf{u}(Q_B(\beta)) > 0, \\ \beta > \beta_{\text{DemParity}}^+ & \text{ if } g_A \mathbf{u}(Q_A^+(\beta)) + g_B \mathbf{u}(Q_B^+(\beta)) < 0. \end{aligned}$$

Proof. Note that we can write

$$\mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta))) = g_A \langle \mathbf{u}, \pi_A \circ r_{\pi_A}^{-1}(\beta) \rangle + g_B \langle \mathbf{u}, \pi_B \circ r_{\pi_B}^{-1}(\beta) \rangle .$$

Since $\mathbf{u}(x)$ is non-decreasing in x , Proposition 5.3 implies that $\beta \mapsto \mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta)))$ is concave in β . Hence, all optimal selection rates β^* lie in an interval $[\beta^-, \beta^+]$. To further characterize this interval, let us compute left- and right-derivatives.

$$\begin{aligned} \partial_+ \mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta))) &= \partial_+ g_A \langle \mathbf{u}, \pi_A \circ r_{\pi_A}^{-1}(\beta) \rangle + \partial_+ g_B \langle \mathbf{u}, \pi_B \circ r_{\pi_B}^{-1}(\beta) \rangle \\ &= g_A \langle \mathbf{u}, \partial_+ (\pi_A \circ r_{\pi_A}^{-1}(\beta)) \rangle + g_B \langle \mathbf{u}, \partial_+ (\pi_B \circ r_{\pi_B}^{-1}(\beta)) \rangle \\ &\stackrel{\text{Lemma 5.3}}{=} g_A \langle \mathbf{u}, \mathbf{e}_{Q_A(\beta)} \rangle + g_B \langle \mathbf{u}, \mathbf{e}_{Q_B(\beta)} \rangle \\ &= g_A \mathbf{u}(Q_A(\beta)) + g_B \mathbf{u}(Q_B(\beta)) . \end{aligned}$$

The same argument shows that

$$\partial_- \mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta))) = g_A \mathbf{u}(Q_A^+(\beta)) + g_B \mathbf{u}(Q_B^+(\beta)) .$$

By concavity of $\mathcal{U}((r_{\pi_A}^{-1}(\beta), r_{\pi_B}^{-1}(\beta)))$, a positive right derivative at β implies that $\beta < \beta^*$ for all β^* satisfying (17), and similarly, a negative left derivative at β implies that $\beta > \beta^*$ for all β^* satisfying (17). $\square$

With a result of the above form, we can now easily prove statements such as that in Corollary 3.3 (see appendix C for proofs), by fixing a selection rate of interest (e.g. β_0) and inverting the

inequalities in Theorem 6.1 to find the exact population proportions under which, for example, DemParity results in a higher selection rate than β_0 .

6.2 EqOpt and General Constraints

Next, we will provide a theorem that gives an explicit characterization for the range of selection rates β_A for A when the bank loans according to EqOpt. Observe that the EqOpt objective corresponds to solving the following linear program:

$$\max_{\tau=(\tau_A, \tau_B) \in [0,1]^{2C}} \mathcal{U}(\tau) \quad \text{s.t.} \quad \langle \mathbf{w}_A \circ \pi_A, \tau_A \rangle = \langle \mathbf{w}_B \circ \pi_B, \tau_B \rangle, \quad (18)$$

where $\mathbf{w}_j = \frac{\rho}{\langle \rho, \pi_j \rangle}$. This problem is similar to the demographic parity optimization in (17), except for the fact that the constraint includes the weights. Whereas we parameterized demographic parity solutions in terms of the acceptance rate β in equation (17), we will parameterize equation (18) in terms of the true positive rate (TPR), $t := \langle \mathbf{w}_A \circ \pi_A, \tau_A \rangle$. Thus, (18) becomes

$$\max_{t \in [0, t_{\max}]} \max_{(\tau_A, \tau_B) \in [0,1]^{2C}} \sum_{j \in \{A, B\}} g_j \mathcal{U}_j(\tau_j) \quad \text{s.t.} \quad \langle \mathbf{w}_j \circ \pi_j, \tau_j \rangle = t, \quad j \in \{A, B\}, \quad (19)$$

where $t_{\max} = \min_{j \in \{A, B\}} \{\langle \pi_j, \mathbf{w}_j \rangle\}$ is the largest possible TPR. The magenta EO curve in Figure 3 illustrates that feasible solutions to this optimization problem lie on a curve parametrized by t . Note that the objective function decouples for $j \in \{A, B\}$ for the inner optimization problem,

$$\max_{\tau_j \in [0,1]^C} \sum_{j \in \{A, B\}} g_j \mathcal{U}_j(\tau_j) \quad \text{s.t.} \quad \langle \mathbf{w}_j \circ \pi_j, \tau_j \rangle = t. \quad (20)$$

We will now show that all optimal solutions for this inner optimization problem are π_j -a.e. equal to a policy in $\mathcal{T}_{\text{thresh}}(\pi_j)$, and thus can be written as $r_{\pi_j}^{-1}(\beta_j)$, depending only on the resulting selection rate.

Proof of Proposition 5.2 for EqOpt. We apply Lemma 5.2 to the inner optimization in (20) with $\mathbf{v}(x) = \mathbf{u}(x)$ and $\mathbf{w}(x) = \frac{\rho(x)}{\langle \rho, \pi_j \rangle}$. The claim follows from the assumption that $\mathbf{u}(x)/\rho(x)$ is increasing by optimizing over t . $\square$

This selection rate β_j is uniquely determined by the TPR t (proof appears in Appendix B.1):

Lemma 6.1. *Suppose that $\mathbf{w}(x) > 0$ for all x . Then the function*

$$T_{j, \mathbf{w}_j}(\beta) := \langle r_{\pi_j}^{-1}(\beta), \pi_j \circ \mathbf{w}_j \rangle$$

is a bijection from $[0, 1]$ to $[0, \langle \pi_j, \mathbf{w} \rangle]$.

Hence, for any $t \in [0, t_{\max}]$, the mapping from TPR to acceptance rate, $T_{j, \mathbf{w}_j}^{-1}(t)$, is well defined and any solution to (20) is π_j -a.e. equal to the policy $r_{\pi_j}^{-1}(T_{j, \mathbf{w}_j}^{-1}(t))$. Thus (19) reduces to

$$\max_{t \in [0, t_{\max}]} \sum_{j \in \{A, B\}} g_j \mathcal{U}_j \left(r_{\pi_j}^{-1} \left(T_{j, \mathbf{w}_j}^{-1}(t) \right) \right). \quad (21)$$

The above expression parametrizes the optimization problem in terms of a single variable. We shall show that the above expression is in fact a *concave* function in t , and hence the set of optimal selection rates can be characterized by first order conditions. This is presented formally in the following theorem:

Theorem 6.2 (Selection rates for Eq0pt). *The set of optimal selection rates β^* for group A satisfying (19) forms a continuous interval $[\beta_{\text{Eq0pt}}^-, \beta_{\text{Eq0pt}}^+]$, such that for any $\beta \in [0, 1]$, we have*

$$\begin{aligned} \beta < \beta_{\text{Eq0pt}}^- & \text{ if } g_A \frac{u(Q_A(\beta))}{w_A(Q_A(\beta))} + g_B \frac{u(Q_B(G_w^{(A \rightarrow B)}(\beta)))}{w_B(Q_B(G_w^{(A \rightarrow B)}(\beta)))} > 0, \\ \beta > \beta_{\text{Eq0pt}}^+ & \text{ if } g_A \frac{u(Q_A^+(\beta))}{w_A(Q_A^+(\beta))} + g_B \frac{u(Q_B^+(G_w^{(A \rightarrow B)}(\beta)))}{w_B(Q_B^+(G_w^{(A \rightarrow B)}(\beta)))} < 0. \end{aligned}$$

Here, $G_w^{(A \rightarrow B)}(\beta) := T_{B, w_B}^{-1}(T_{A, w_A}^{-1}(\beta))$ denotes the (well-defined) map from selection rates β_A for A to the selection rate β_B for B such that the policies $\tau_A^* := r_{\pi_A}^{-1}(\beta_A)$ and $\tau_B^* := r_{\pi_B}^{-1}(\beta_B)$ satisfy the constraint in (18).

Proof. Starting with the equivalent problem in (21), we use the concavity result of Lemma B.1. Because the objective function is the positive weighted sum of two concave functions, it is also concave. Hence, all optimal true positive rates t^* lie in an interval $[t^-, t^+]$. To further characterize $[t^-, t^+]$, we can compute left- and right-derivatives, again using the result of Lemma B.1.

$$\begin{aligned} \partial_+ \sum_{j \in \{A, B\}} g_j \mathcal{U}_j \left(r_{\pi_j}^{-1}(T_{j, w_j}^{-1}(t)) \right) &= g_A \partial_+ \mathcal{U}_A \left(r_{\pi_A}^{-1}(T_{A, w_A}^{-1}(t)) \right) + g_B \partial_+ \mathcal{U}_B \left(r_{\pi_B}^{-1}(T_{B, w_B}^{-1}(t)) \right) \\ &= g_A \frac{u(Q_A(T_{A, w_A}^{-1}(t)))}{w_A(Q_A(T_{A, w_A}^{-1}(t)))} + g_B \frac{u(Q_B(T_{B, w_B}^{-1}(t)))}{w_B(Q_B(T_{B, w_B}^{-1}(t)))} \end{aligned}$$

The same argument shows that

$$\partial_- \sum_{j \in \{A, B\}} g_j \mathcal{U}_j \left(r_{\pi_j}^{-1}(T_{j, w_j}^{-1}(t)) \right) = g_A \frac{u(Q_A^+(T_{A, w_A}^{-1}(t)))}{w_A(Q_A^+(T_{A, w_A}^{-1}(t)))} + g_B \frac{u(Q_B^+(T_{B, w_B}^{-1}(t)))}{w_B(Q_B^+(T_{B, w_B}^{-1}(t)))}.$$

By concavity, a positive right derivative at t implies that $t < t^*$ for all t^* satisfying (21), and similarly, a negative left derivative at t implies that $t > t^*$ for all t^* satisfying (21).

Finally, by Lemma 6.1, this interval in t uniquely characterizes an interval of acceptance rates. Thus we translate directly into a statement about the selection rates β for group A by seeing that $T_{A, w_A}^{-1}(t) = \beta$ and $T_{B, w_B}^{-1}(t) = G_w^{(A \rightarrow B)}(\beta)$. $\square$

Lastly, we remark that the results derived in this section go through verbatim for any linear constraint of the form $\langle w, \pi_A \circ \tau_A \rangle = \langle w, \pi_B \circ \tau_B \rangle$, as long as $u(x)/w(x)$ is increasing in x , and $w(x) > 0$.




Figure 4: The empirical payback rates as a function of credit score and CDF for both groups from the TransUnion TransRisk dataset.

7 Simulations

We examine the outcomes induced by fairness constraints in the context of FICO scores for two race groups. FICO scores are a proprietary classifier widely used in the United States to predict credit worthiness. Our FICO data is based on a sample of 301,536 TransUnion TransRisk scores from 2003 [US Federal Reserve, 2007], preprocessed by Hardt et al. [2016]. These scores, corresponding to x in our model, range from 300 to 850 and are meant to predict credit risk. Empirical data labeled by race allows us to estimate the distributions π_j , where j represents race, which is restricted to two values: white non-Hispanic (labeled “white” in figures), and black. Using national demographic data, we set the population proportions to be 18% and 82%.

Individuals were labeled as defaulted if they failed to pay a debt for at least 90 days on at least one account in the ensuing 18-24 month period; we use this data to estimate the success probability given score, $\rho_j(x)$, which we allow to vary by group to match the empirical data (see Figure 4). Our outcome curve framework allows for this relaxation; however, this discrepancy can also be attributed to group-dependent mismeasurement of score, and adjusting the scores accordingly would allow for a single $\rho(x)$. We use the success probabilities to define the affine utility and score change functions defined in Example 2.1. We model individual penalties as a score drop of $c_- = -150$ in the case of a default, and an increase of $c_+ = 75$ in the case of successful repayment.

In Figure 5, we display the empirical CDFs along with selection rates resulting from different loaning strategies for two different settings of bank utilities. In the case that the bank experiences a loss/profit ratio of $\frac{u_-}{u_+} = -10$, no fairness criteria surpass the active harm rate β_0 ; however, in the case of $\frac{u_-}{u_+} = -4$, **DemParity** overloans, in line with the statement in Corollary 3.3.

These results are further examined in Figure 6, which displays the normalized outcome curves and the utility curves for both the white and the black group. To plot the **MaxUtil** utility curves, the group that is not on display has selection rate fixed at β^{MaxUtil} . In this figure, the top panel corresponds to the average change in credit scores for each group under different loaning rates β ; the bottom panels shows the corresponding *total* utility $\mathcal{U}$ (summed over both groups and weighted by group population sizes) for the bank.




Figure 5: The empirical CDFs of both groups are plotted along with the decision thresholds resulting from **MaxUtil**, **DemParity**, and **EqOpt** for a model with bank utilities set to (a) $\frac{u_-}{u_+} = -4$ and (b) $\frac{u_-}{u_+} = -10$. The threshold for active harm is displayed; in (a) **DemParity** causes active harm while in (b) it does not. **EqOpt** and **MaxUtil** never cause active harm.

Figure 6 highlights that the position of the utility optima in the lower panel determines the loan (selection) rates. In this specific instance, the utility and change ratios are fairly close, $\frac{u_-}{u_+} = -4$, and $\frac{c_-}{c_+} = -2$, meaning that the bank’s profit motivations align with individual outcomes to some extent. Here, we can see that **EqOpt** loans much closer to optimal than **DemParity**, similar to the setting suggested by Corollary 3.2.

Although one might hope for decisions made under fairness constraints to positively affect the black group, we observe the opposite behavior. The **MaxUtil** policy (solid orange line) and the **EqOpt** policy result in similar expected credit score change for the black group. However, **DemParity** (dashed green line) causes a negative expected credit score change in the black group, corresponding to active harm. For the white group, the bank utility curve has almost the same shape under the fairness criteria as it does under **MaxUtil**, the main difference being that fairness criteria lowers the total expected profit from this group.

This behavior stems from a discrepancy in the outcome and profit curves for each population. While incentives for the bank and positive results for individuals are somewhat aligned for the majority group, under fairness constraints, they are more heavily misaligned in the minority group, as seen in graphs (left) in Figure 6. We remark that in other settings where the *unconstrained* profit maximization is misaligned with individual outcomes (e.g., when $\frac{u_-}{u_+} = -10$), fairness criteria may perform more favorably for the minority group by pulling the utility curve into a shape consistent with the outcome curve.

By analyzing the resulting affects of **MaxUtil**, **DemParity**, and **EqOpt** on actual credit score lending data, we show the applicability of our model to real-world applications. In particular, some results shown in Section 3 hold empirically for the FICO TransUnion TransRisk scores.




Figure 6: The outcome and utility curves are plotted for both groups against the group selection rates. The relative positions of the utility maxima determine the position of the decision rule thresholds. We hold $\frac{u_-}{u_+} = -4$ as fixed.

8 Conclusion and Future Work

We argue that without a careful model of delayed outcomes, we cannot foresee the impact a fairness criterion would have if enforced as a constraint on a classification system. However, if such an accurate outcome model is available, we show that there are more direct ways to optimize for positive outcomes than via existing fairness criteria.

Our formal framework exposes a concise, yet expressive way to model outcomes via the expected change in a variable of interest caused by an institutional decision. This leads to the natural concept of an outcome curve that allows us to interpret and compare solutions effectively. In essence, the formalism we propose requires us to understand the two-variable causal mechanism that translates decisions to outcomes. Depending on the application, such an understanding might necessitate greater domain knowledge and additional research into the specifics of the application. This is consistent with much scholarship that points to the context-sensitive nature of fairness in machine learning.

An interesting direction for future work is to consider other characteristics of impact beyond the change in population *mean*. Variance and individual-level outcomes are natural and important considerations. Moreover, it would be interesting to understand the robustness of outcome optimization to modeling and measurement errors.

Acknowledgements

We thank Lily Hu, Aaron Roth, and Cathy O’Neil for discussions and feedback on an earlier version of the manuscript. We thank the students of CS294: Fairness in Machine Learning (Fall 2017, University of California, Berkeley) for inspiring class discussions and comments on a presentation that was a precursor of this work. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814.

References

- Solon Barocas and Andrew D. Selbst. Big data’s disparate impact. *California Law Review*, 104, 2016.
- Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building classifiers with independency constraints. In *Proc. IEEE ICDMW*, ICDMW ’09, pages 13–18, 2009.
- Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *FATML*, 2016.
- Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. *arXiv preprint arXiv:1706.09847*, 2017.
- Executive Office of the President. Big data: A report on algorithmic systems, opportunity, and civil rights. Technical report, White House, May 2016.
- Dean P Foster and Rakesh V Vohra. An economic argument for affirmative action. *Rationality and Society*, 4(2):176–188, 1992.
- Andreas Fuster, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. Predictably unequal? the effects of machine learning on credit markets. *SSRN*, 2017.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proc. 30th NIPS*, 2016.
- Lily Hu and Yiling Chen. A short-term intervention for long-term fairness in the labor market. In *Proc. 27th WWW*, 2018.
- Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. In *Proc. 30th NIPS*, pages 325–333, 2016.
- Alexandra Kaley, Frank Dobbin, and Erin Kelly. Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies. *American Sociological Review*, 71(4):589–617, 2006.
- Stephen N. Keith, Robert M. Bell, August G. Swanson, and Albert P. Williams. Effects of affirmative action in medical schools. *New England Journal of Medicine*, 313(24):1519–1525, 1985.
- Niki Kilbertus, Mateo Rojas-Carulla, Giambattista Parascandolo, Moritz Hardt, Dominik Janzing, and Bernhard Schölkopf. Avoiding discrimination through causal reasoning. In *In Proc. 30th NIPS*, pages 656–666, 2017.
- Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *Proc. 8th ITCS*, 2017.
- Matt J. Kusner, Joshua R. Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *In Proc. 30th NIPS*, pages 4069–4079, 2017.
- Razieh Nabi and Ilya Shpitser. Fair inference on outcomes. *arXiv:1705.10378v1*, 2017.

Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. In *Advances in Neural Information Processing Systems 30*, pages 5684–5693, 2017.

Stephen Ross and John Yinger. *The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement*. MIT Press, Cambridge, 2006.

US Federal Reserve. Report to the congress on credit scoring and its effects on the availability and affordability of credit, 2007.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P. Gummadi. Fairness Constraints: Mechanisms for Fair Classification. In *Proc. 20th AISTATS*, pages 962–970. PMLR, 2017.

A Optimality of Threshold Policies

A.1 Proof of Lemma 5.1

We begin with the first statement of the lemma. Suppose $\tau_j \cong_{\pi_j} \tau'_j$. Then there exists a set $\mathcal{S} \subset \mathcal{X}$ such that $\pi_j(x) = 0$ for all $x \in \mathcal{S}$, and for all $x \notin \mathcal{S}$, $\tau_j(x) = \tau'_j(x)$. Thus,

$$\begin{aligned} r_{\pi}(\tau_j) - r_{\pi_j}(\tau'_j) &= \sum_{x \in \mathcal{X}} \pi_j(x)(\tau_j(x) - \tau'_j(x)) \\ &= \sum_{x \in \mathcal{S}} \pi_j(x)(\tau_j(x) - \tau'_j(x)) = 0. \end{aligned}$$

Conversely, suppose that $r_{\pi_j}(\tau_j) = r_{\pi_j}(\tau'_j)$. Let $\tau_j = \tau_{c,\gamma}$ and $\tau'_j = \tau_{c',\gamma'}$ as in Definition 5.1. We now have the following cases:

1. Case 1: $c = c'$. Then $\tau_j(x) = \tau'_j(x)$ for all $x \in \mathcal{X} - \{c\}$. Hence,

$$0 = r_{\pi}(\tau_j) - r_{\pi_j}(\tau'_j) = \pi(x)(\tau_j(c) - \tau'_j(c)).$$

This implies that either $\tau_j(c) = \tau'_j(c)$, and thus $\tau_j(x) = \tau'_j(x)$ for all $x \in \mathcal{X}$, or otherwise $\pi(c) = 0$, in which case we still have $\tau_j \cong_{\pi_j} \tau'_j$ (since the two policies agree every outside the set $\{c\}$).

2. Case 2: $c \neq c'$. We assume without loss of generality that $c' < c \leq C$. Since the policies $\tau_{c',1}$ and $\tau_{c'+1,0}$ are identity for $c' < C$, we may also assume without loss of generality that $\gamma' \in [0, 1)$. Thus for all $x \in \mathcal{S} := \{c', c' + 1, \dots, C\}$, we have $\tau'_j(x) < \tau_j(x)$. This implies that

$$\begin{aligned} 0 &= r_{\pi}(\tau_j) - r_{\pi_j}(\tau'_j) \\ &= \sum_{x \in \mathcal{S}} \pi_j(x)(\tau_j(x) - \tau'_j(x)) \\ &\geq \min_{x \in \mathcal{S}} (\tau_j(c) - \tau'_j(x)) \cdot \sum_{x \in \mathcal{S}} \pi(x). \end{aligned}$$

Since $\min_{x \in \mathcal{S}} (\tau_j(c) - \tau'_j(x)) > 0$, it follows that $\sum_{x \in \mathcal{S}} \pi_j(x) = 0$, whence $\tau_j \cong_{\pi_j} \tau'_j$.

Next, we show that r_π is a bijection from $\mathcal{T}_{\text{thresh}}(\pi) \rightarrow [0, 1]$. That r_π is injective follows immediately from the fact if $r_{\pi_j}(\tau) = r_{\pi_j}(\tau'_j)$, then $\tau_j \cong_{\pi_j} \tau'_j$. To show it is surjective, we exhibit for every $\beta \in [0, 1]$ a threshold policy $\tau_{c,\gamma}$ for which $r_{\pi_j}(\tau_{c,\gamma}) = \beta$. We may assume $\beta < 1$, since the all-ones policy has a selection rate of 1.

Recall the definition of the inverse CDF

$$Q_j(\beta) := \operatorname{argmax}\{c : \sum_{x=c}^C \pi(x) > \beta\}.$$

Since $\beta < 1$, $Q_j(\beta) \leq C$. Let $\beta_+ = \sum_{x=Q_j(\beta)}^C \pi(x)$, and let $\beta_- = \sum_{x=Q_j(\beta)+1}^C \pi(x)$. Note that by definition, we have $\beta_- \leq \beta < \beta_+$, and $\beta_+ - \beta_- = \pi(Q_j(\beta))$. Hence, if we define $\gamma = \frac{\beta - \beta_-}{\beta_+ - \beta_-}$, we have

$$r_{\pi_j}(\tau_{Q_j(\beta),\gamma}) = \pi(Q_j(\beta))\gamma + \sum_{x=Q_j(\beta)+1}^C \pi(x) = \beta_- + (\beta_+ - \beta_-)\gamma = \beta_- + \beta - \beta_- = \beta.$$

A.2 Proof of Lemma 5.2

Given $\tau \in [0, 1]^C$, we define the *normal cone* at τ as $\text{NC}(\tau) := \text{ConicalHull}\{\mathbf{z} : \tau + \mathbf{z} \in [0, 1]^C\}$. We can describe $\text{NC}(\tau)$ explicitly as:

$$\text{NC}(\tau) := \{\mathbf{z} \in \mathbb{R}^C : z_i \leq 0 \text{ if } \tau_i = 0, z_i \geq 0 \text{ if } \tau_i = 1\}.$$

Immediately from the above definition, we have the following useful identity, which is that for any vector $\mathbf{g} \in \mathbb{R}^C$,

$$\langle \mathbf{g}, \mathbf{z} \rangle \leq 0 \quad \forall \mathbf{z} \in \text{NC}(\tau), \quad \text{if and only if} \quad \forall x \in \mathcal{X}, \begin{cases} \tau(x) = 0 & \mathbf{g}(x) < 0 \\ \tau(x) = 1 & \mathbf{g}(x) > 0 \\ \tau(x) \in [0, 1] & \mathbf{g}(x) = 0 \end{cases}. \quad (22)$$

Now consider the optimization problem (12). By the first order KKT conditions, we know that for any optimizer τ_* of the above objective, there exists some $\hat{\lambda} \in \mathbb{R}$ such that, for all $\mathbf{z} \in \text{NC}(\tau_*)$

$$\langle \mathbf{z}, \mathbf{v} \circ \pi + \hat{\lambda} \pi \circ \mathbf{w} \rangle \leq 0.$$

By (22), we must have that

$$\tau_*(x) = \begin{cases} 0 & \pi(x)(\mathbf{v}(x) + \hat{\lambda} \mathbf{w}(x)) < 0 \\ 1 & \pi(x)(\mathbf{v}(x) + \hat{\lambda} \mathbf{w}(x)) > 0 \\ \in [0, 1] & \pi(x)(\mathbf{v}(x) + \hat{\lambda} \mathbf{w}(x)) = 0 \end{cases}.$$

Now $\tau_*(x)$ is not necessarily a threshold policy. To conclude the theorem, it suffices to exhibit a threshold policy $\tilde{\tau}_*$ such that $\tau_*(x) \cong_{\pi} \tilde{\tau}_*$. (Note that $\tilde{\tau}_*(x)$ will also be feasible for the constraint, and have the same objective value; hence $\tilde{\tau}_*$ will be optimal as well.)

Given τ_* and $\hat{\lambda}$, let $c_* = \min\{c \in \mathcal{X} : v(x) + \hat{\lambda}w(x) \geq 0\}$. If either (a) $w(x) = 0$ for all $x \in \mathcal{X}$ and $v(x)$ is strictly increasing or (b) $v(x)/w(x)$ is strictly increasing, then the modified policy

$$\tilde{\tau}_*(x) = \begin{cases} 0 & x < c_* \\ \tau_*(x) & x = c_* \\ 1 & x > c_* \end{cases},$$

is a threshold policy, and $\tau_*(x) \cong_{\pi} \tilde{\tau}_*$. Moreover, $\langle w, \tilde{\tau}_* \rangle = \langle w, \tau_* \rangle$ and $\langle \pi, \tilde{\tau}_* \rangle = \langle \pi, \tau_* \rangle$, which implies that $\tilde{\tau}_*$ is an optimal policy for the objective in Lemma 5.2.

A.3 Proof of Lemma 5.3

We shall prove

$$\partial_+ \left(\pi_j \circ r_{\pi_j}^{-1}(\beta) \right) = e_{Q_j(\beta)}, \quad (23)$$

where the derivative is with respect to β . The computation of the left-derivative is analogous. Since we are concerned with right-derivatives, we shall take $\beta \in [0, 1)$. Since $\pi_j \circ r_{\pi_j}^{-1}(\beta)$ does not depend on the choice of representative for $r_{\pi_j}^{-1}$, we can choose a canonical representation for $r_{\pi_j}^{-1}$. In Section A.1, we saw that the threshold policy $\tau_{Q_j(\beta), \gamma(\beta)}$ had acceptance rate β , where we had defined

$$\beta_+ = \sum_{x=Q_j(\beta)}^C \pi(x) \quad \text{and} \quad \beta_- = \sum_{x=Q_j(\beta)+1}^C \pi(x), \quad (24)$$

$$\gamma(\beta) = \frac{\beta - \beta_-}{\beta_+ - \beta_-}. \quad (25)$$

Note then that for each x , $\tau_{Q_j(\beta), \gamma(\beta)}(x)$ is piece-wise linear, and thus admits left and right derivatives. We first claim that

$$\forall x \in \mathcal{X} \setminus \{Q_j(\beta)\}, \quad \partial_+ \tau_{Q_j(\beta), \gamma(\beta)}(x) = 0. \quad (26)$$

To see this, note that $Q_j(\beta)$ is right continuous, so for all ϵ sufficiently small, $Q_j(\beta + \epsilon) = Q_j(\beta)$. Hence, for all ϵ sufficiently small and all $x \neq Q_j(\beta)$, we have $\tau_{Q_j(\beta+\epsilon), \gamma(\beta+\epsilon)}(x) = \tau_{Q_j(\beta), \gamma(\beta)}(x)$, as needed. Thus, Equation (26) implies that $\partial_+ \pi_j \circ r_{\pi_j}^{-1}(\beta)$ is supported on $x = Q_j(\beta)$, and hence

$$\partial_+ \pi_j \circ r_{\pi_j}^{-1}(\beta) = \partial_+ \pi_j(x) \tau_{Q_j(\beta), \gamma(\beta)}(x) \Big|_{x=Q_j(\beta)} \cdot e_{Q_j(\beta)}.$$

To conclude, we must show that $\partial_+ \pi_j(x) \tau_{Q_j(\beta), \gamma(\beta)}(x) \Big|_{x=Q_j(\beta)} = 1$. To show this, we have

$$\begin{aligned} 1 &= \partial_+(\beta) \\ &= \partial_+(r_{\pi_j}(\tau_{Q_j(\beta), \gamma(\beta)})) \quad \text{since} \quad r_{\pi_j}(\tau_{Q_j(\beta), \gamma(\beta)}) = \beta \quad \forall \beta \in [0, 1) \\ &= \partial_+ \left(\sum_{x \in \mathcal{X}} \pi(x) \cdot \tau_{Q_j(\beta), \gamma(\beta)}(x) \right) \end{aligned}$$

$$= \partial_+ \pi(x) \cdot \tau_{Q_j(\beta), \gamma(\beta)}(x) \Big|_{x=Q_j(\beta)}, \text{ as needed.}$$

B Characterization of Fairness Solutions

B.1 Derivative Computation for EqOpt

In this section, we prove Lemma 6.1, which we recall below.

Lemma 6.1. *Suppose that $w(x) > 0$ for all x . Then the function*

$$T_{j, w_j}(\beta) := \langle r_{\pi_j}^{-1}(\beta), \pi_j \circ w_j \rangle$$

is a bijection from $[0, 1]$ to $[0, \langle \pi_j, w \rangle]$.

We will prove Lemma 6.1 in tandem with the following derivative computation which we applied in the proof of Theorem 6.2.

Lemma B.1. *The function*

$$\mathcal{U}_j(t; w_j) := \mathcal{U}_j \left(r_{\pi_j}^{-1} \left(T_{j, w_j}^{-1}(t) \right) \right)$$

is concave in t and has left and right derivatives

$$\partial_+ \mathcal{U}_j(t; w_j) = \frac{u(Q_j(T_{j, w_j}^{-1}(t)))}{w_j(Q_j(T_{j, w_j}^{-1}(t)))} \quad \text{and} \quad \partial_- \mathcal{U}_j(t; w_j) = \frac{u(Q_j^+(T_{j, w_j}^{-1}(t)))}{w_j(Q_j^+(T_{j, w_j}^{-1}(t)))}.$$

Proof of Lemmas 6.1 and B.1. Consider a $\beta \in [0, 1]$. Then, $\pi_j \circ r_{\pi_j}^{-1}(\beta)$ is continuous and left and right differentiable by Lemma 5.3, and its left and right derivatives are indicator vectors $e_{Q_j(\beta)}$ and $e_{Q_j^+(\beta)}$, respectively. Consequently, $\beta \mapsto \langle w_j, \pi_j \circ r_{\pi_j}^{-1}(\beta) \rangle$ has left and right derivatives $w_j(Q(\beta))$ and $w_j(Q^+(\beta))$, respectively; both of which are both strictly positive by the assumption $w(x) > 0$. Hence, $T_{j, w_j}(\beta) = \langle w_j, \pi_j \circ r_{\pi_j}^{-1}(\beta) \rangle$ is strictly increasing in β , and so the map is injective. It is also surjective because $\beta = 0$ induces the policy $\tau_j = \mathbf{0}$ and $\beta = 1$ induces the policy $\tau_j = \mathbf{1}$ (up to π_j -measure zero). Hence, $T_{j, w_j}(\beta)$ is an order preserving bijection with left- and right-derivatives, and we can compute the left and right derivatives of its inverse as follows:

$$\partial_+ T_{j, w_j}^{-1}(t) = \frac{1}{\partial_+ T_{j, w_j}(\beta) \Big|_{\beta=T_{j, w_j}^{-1}(t)}} = \frac{1}{w_j(Q_j(T_{j, w_j}^{-1}(t)))},$$

and similarly, $\partial_- T_{j, w_j}^{-1}(t) = \frac{1}{w_j(Q_j^+(T_{j, w_j}^{-1}(t)))}$. Then we can compute that

$$\begin{aligned} \partial_+ \mathcal{U}_j(r_{\pi_j}(T_{j, w_j}^{-1}(t))) &= \partial_+ \mathcal{U}(r_{\pi_j}(\beta)) \Big|_{\beta=T_{j, w_j}^{-1}(t)} \cdot \partial_+ T_{j, w_j}(\sup(t)) \\ &= \frac{u(Q_j(T_{j, w_j}^{-1}(t)))}{w_j(Q_j(T_{j, w_j}^{-1}(t)))}. \end{aligned}$$

and similarly $\partial_- \mathcal{U}_j(r_{\pi_j}(T_{j, w_j}^{-1}(t))) = \frac{u(Q_j^+(T_{j, w_j}^{-1}(t)))}{w_j(Q_j^+(T_{j, w_j}^{-1}(t)))}$. One can verify that for all $t_1 < t_2$, one has that

$\partial_+ \mathcal{U}_j(r_{\pi_j}(T_{j,w_j}^{-1}(t_1))) \geq \partial_- \mathcal{U}_j(r_{\pi_j}(T_{j,w_j}^{-1}(t_2)))$, and that for all t , $\partial_+ \mathcal{U}_j(r_{\pi_j}(T_{j,w_j}^{-1}(t))) \leq \partial_- \mathcal{U}_j(r_{\pi_j}(T_{j,w_j}^{-1}(t)))$. These facts establish that the mapping $t \mapsto \mathcal{U}_j(r_{\pi_j}(T_{j,w_j}^{-1}(t)))$ is concave. $\square$

B.2 Characterizations Under Soft Constraints

Given a convex penalty $\Phi : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, and $\lambda \in \mathbb{R}_{\geq 0}$, one can write down the general form for soft constrained utility optimization

$$\max_{\tau=(\tau_A, \tau_B)} \mathcal{U}(\tau) - \lambda \Phi(\langle \mathbf{w}_A \circ \pi_A, \tau_A \rangle - \langle \mathbf{w}_B \circ \pi_B, \tau_B \rangle), \quad (27)$$

where $\mathbf{w}_A$ and $\mathbf{w}_B$ represent generic constraints. Again, we shall assume that for $j \in \{A, B\}$, $\mathbf{u}(x)/\mathbf{w}_j(x)$ is non-decreasing. Recall that for $\mathbf{w}_j = (1, 1, \dots, 1)$, one recovers the soft version of DemParity, whereas for $\mathbf{w}_j = \frac{\rho}{\langle \rho, \pi_j \rangle}$, one recovers the soft constrained version of Eq0pt.

The same argument presented in Section 6.2 shows that the optimal policies are of the form

$$\tau_j = r_{\pi_j}^{-1}(T_{j,w_j}^{-1}(t_j)),$$

where (t_A, t_B) are solutions to the following optimization problem:

$$\max_{t_A \in [0, \langle \pi_A, \mathbf{w}_A \rangle], t_B \in [0, \langle \pi_B, \mathbf{w}_B \rangle]} g_A \mathcal{U}_A(r_{\pi_A}^{-1}(T_{A,w_A}^{-1}(t_A))) + g_B \mathcal{U}_B(r_{\pi_B}^{-1}(T_{B,w_B}^{-1}(t_B))) - \lambda \Phi(t_A - t_B).$$

The following lemma gives us a first order characterization of these optimal TPRs, (t_A, t_B) .

Lemma B.2. *All optimal policies are equivalent to threshold policies with selection rate (β_A, β_B) which satisfy*

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \left[\begin{bmatrix} \frac{\mathbf{u}(Q_A(\beta_A))}{\mathbf{w}_A(Q_A(\beta_B))} - \lambda \partial_+ \Phi(\Delta), \frac{\mathbf{u}(Q_A^+(\beta_A))}{\mathbf{w}_A(Q_A^+(\beta_A))} - \lambda \partial_- \Phi(\Delta) \\ \frac{\mathbf{u}(Q_B(\beta_B))}{\mathbf{w}_B(Q_B(\beta_B))} + \lambda \partial_- \Phi(\Delta), \frac{\mathbf{u}(Q_B^+(\beta_B))}{\mathbf{w}_B(Q_B^+(\beta_B))} + \lambda \partial_+ \Phi(\Delta) \end{bmatrix} \right], \quad (28)$$

where $\Delta = t_A - t_B = T_{A,w_A}(\beta_A) - T_{B,w_B}(\beta_B)$.

Proof. Let $\partial(\cdot)$ denote the super-gradient set of a concave function. Note that if F is left-and-right differentiable and concave, then $\partial F(x) = [\partial_+ F(x), \partial_- F(x)]$. By concavity of $\mathcal{U}_j$ and convexity of Φ , we must have that

$$\begin{aligned} \begin{bmatrix} 0 \\ 0 \end{bmatrix} &\in \partial \sum_{j \in \{A, B\}} \mathcal{U}_j(r_{\pi_j}^{-1}(T_{j,w_j}^{-1}(t_j))) - \lambda \Phi(t_A - t_B) \\ &= \left[\partial \mathcal{U}_A(r_{\pi_A}^{-1}(T_{A,w_A}^{-1}(t_A))) + \partial_{t_A} \{-\lambda \Phi(t_A - t_B)\} \right. \\ &\quad \left. \partial \mathcal{U}_B(r_{\pi_B}^{-1}(T_{B,w_B}^{-1}(t_B))) + \partial_{t_B} \{-\lambda \Phi(t_A - t_B)\} \right] \\ &= \left[\partial \mathcal{U}_A(r_{\pi_A}^{-1}(T_{A,w_A}^{-1}(t_A))) - \lambda \partial \Phi(t) \Big|_{t=t_A-t_B} \right. \\ &\quad \left. \partial \mathcal{U}_B(r_{\pi_B}^{-1}(T_{B,w_B}^{-1}(t_B))) + \lambda \partial \Phi(t) \Big|_{t=t_A-t_B} \right] \\ &= \left[\left[\partial_+ \mathcal{U}_A(r_{\pi_A}^{-1}(T_{A,w_A}^{-1}(t_A))) - \lambda \partial_+ \Phi(t) \Big|_{t=t_A-t_B}, \partial_- \mathcal{U}_A(r_{\pi_A}^{-1}(T_{A,w_A}^{-1}(t_A))) - \lambda \partial_- \Phi(t) \Big|_{t=t_A-t_B} \right] \right. \\ &\quad \left. \left[\partial_+ \mathcal{U}_B(r_{\pi_B}^{-1}(T_{B,w_B}^{-1}(t_B))) + \lambda \partial_- \Phi(t) \Big|_{t=t_A-t_B}, \partial_- \mathcal{U}_B(r_{\pi_B}^{-1}(T_{B,w_B}^{-1}(t_B))) + \lambda \partial_+ \Phi(t) \Big|_{t=t_A-t_B} \right] \right] \end{aligned}$$

$$\begin{aligned}
&= \left[\begin{aligned} &\frac{u(Q_A(T_{A,w_A}^{-1}(t_A)))}{w_A(Q_A(T_{A,w_A}^{-1}(t_A)))} - \lambda \partial_+ \Phi(t)|_{t=t_A-t_B}, \frac{u(Q_A^+(T_{A,w_A}^{-1}(t_A)))}{w_A(Q_A^+(T_{A,w_A}^{-1}(t_A)))} - \lambda \partial_- \Phi(t)|_{t=t_A-t_B} \\ &\frac{u(Q_B(T_{B,w_B}^{-1}(t_B)))}{w_B(Q_B(T_{B,w_B}^{-1}(t_B)))} + \lambda \partial_- \Phi(t)|_{t=t_A-t_B}, \frac{u(Q_B^+(T_{B,w_B}^{-1}(t_B)))}{w_B(Q_B^+(T_{B,w_B}^{-1}(t_B)))} + \lambda \partial_+ \Phi(t)|_{t=t_A-t_B} \end{aligned} \right] \\
&= \left[\begin{aligned} &\frac{u(Q_A(\beta_A))}{w_A(Q_A(\beta_A))} - \lambda \partial_+ \Phi(t)|_{t=t_A-t_B}, \frac{u(Q_A^+(\beta_A))}{w_A(Q_A^+(\beta_A))} - \lambda \partial_- \Phi(t)|_{t=t_A-t_B} \\ &\frac{u(Q_B(\beta_B))}{w_B(Q_B(\beta_B))} + \lambda \partial_- \Phi(t)|_{t=t_A-t_B}, \frac{u(Q_B^+(\beta_B))}{w_B(Q_B^+(\beta_B))} + \lambda \partial_+ \Phi(t)|_{t=t_A-t_B} \end{aligned} \right].
\end{aligned}$$

Substituting $\Delta = t_A - t_B = T_{A,w_A}(\beta_A) - T_{B,w_B}(\beta_B)$ concludes the proof. $\square$

In general, a closed form solution for the soft constrained problem may be difficult to state. However, for the case of $\Phi(t) = |t|$, we can state an explicit closed form solution:

Proposition B.1 (Special case of $\Phi(t) = |t|$). *Let $\Phi(t) = |t|$, fix λ , and let $[\beta_A^{\lambda,-}, \beta_A^{\lambda,+}]$ denote the interval of optimal selection rates for Equation (27) with regularization λ . Finally, suppose that for any optimal **MaxUtil** selection rates $(\beta_A^{\text{MaxUtil}}, \beta_B^{\text{MaxUtil}})$, one has $T_{A,w_A}(\beta_A^{\text{MaxUtil}}) < T_{B,w_B}(\beta_B^{\text{MaxUtil}})$. Let $[\beta_A^-, \beta_A^+]$ denote the optimal loan rates in (27). Then there exists a λ_* such that, for $\lambda \geq \lambda_*$, $[\beta_A^-, \beta_A^+]$ coincides with the hard constrained solution. Moreover, for $\lambda < \lambda_*$, any $\beta \in [0, 1]$ satisfies*

$$\begin{aligned}
\beta &< \beta_A^{\lambda,-} \quad \text{if} \quad g_A \frac{u(Q_A(\beta))}{w_A(Q_A(\beta))} + \sigma_* \lambda > 0 \\
\beta &> \beta_A^{\lambda,+} \quad \text{if} \quad g_A \frac{u(Q_A^+(\beta))}{w_A(Q_A^+(\beta))} + \sigma_* \lambda < 0.
\end{aligned}$$

Proof. Given a set of optimal constraint values $(t_A, t_B) = (T_{A,w_A}(\beta_A), T_{B,w_B}(\beta_B))$ for optimal selection rates (β_A, β_B) for a given parameter λ . By Proposition B.2 below, it follows that if $t_A = t_B$ for all optimal solutions, then for all $\lambda' \geq \lambda$, all optimal solutions must also have $t_A = t_B$.

Hence, it suffices to show that (a) there exists a finite λ such that all solutions must have $t_A = t_B$, and (b) if $t_A \neq t_B$, then the display in (B.1) holds.

To prove (a) and (b), suppose $t_A \neq t_B$. By Proposition B.2 below and the fact that $T_{A,w_A}(\beta_A^{\text{MaxUtil}}) < T_{B,w_B}(\beta_B^{\text{MaxUtil}})$, we have $t_A < t_B$. Moreover we can compute that

$$\partial \Phi(t) = \begin{cases} \{1\} & t > 0 \\ [-1, 1] & t = 0 \\ \{-1\} & t < 0 \end{cases}$$

it follows from the first order condition in Lemma B.2 that, if $t_A \neq t_B$

$$0 \in \left[-\frac{u(Q_A^+(\beta_A))}{w_A(Q_A^+(\beta_A))} + \lambda, \frac{u(Q_A(\beta_A))}{w_A(Q_A(\beta_B))} + \lambda \right], \quad (29)$$

which immediately implies point (b). Point (a) follows from the above display by noting that, since $w_j(x) > 0$ and $u(x) < \infty$ for all x , where exists a λ sufficiently large such that (29) cannot hold for any β_A . $\square$

B.3 Qualitative Behavior of Soft Constraints

We now present a proposition which formalizes the intuition that soft constraints interpolate between **MaxUtil** and the general hard constraint (18) in Section 6.2 (for arbitrary $\mathbf{w}$, not just for **EqOpt**). Because optimal policies may not be unique, we define the solution sets

$$\mathbf{P}(\lambda) := \{(\tau_A, \tau_B) : (\tau_A, \tau_B) \text{ solves (27) with parameter } \lambda\},$$

with the set $\mathbf{P}(\infty)$ denoting the set of solutions to (18).

At a high level, we parameterize the soft constrained solution in terms of the value of the constraint $t_A = \langle \tau_A, \mathbf{w}_A \circ \pi_A \rangle$ for **A** and the difference in constraint values $\Delta = \langle \tau_A, \mathbf{w}_A \circ \pi_A \rangle - \langle \tau_B, \mathbf{w}_B \circ \pi_B \rangle$, where $(\tau_A, \tau_B) \in \mathbf{P}(\lambda)$. We show that t_A interpolates between the value of the constraint on **A** at $\lambda = 0$ and at $\lambda = \infty$, and that Δ interpolates between the difference at $\lambda = 0$ (**MaxUtil**) and at $\Delta = 0$ at $\lambda = \infty$. To be rigorous, we note that the possible values for t_A and Δ for each λ are actually contiguous intervals. Hence, to make the interpolation precise, we define the following partial order on such intervals:

Definition B.1 (Interval order). Let $\mathcal{S}_1, \mathcal{S}_2$ be two intervals. We say that $\mathcal{S}_1 \prec \mathcal{S}_2$ if $\max \{x \in \mathcal{S}_1\} < \min \{x \in \mathcal{S}_2\}$ and $\mathcal{S}_1 \preceq \mathcal{S}_2$ if both $\max \{x \in \mathcal{S}_1\} \leq \max \{x \in \mathcal{S}_2\}$ and $\min \{x \in \mathcal{S}_1\} \leq \min \{x \in \mathcal{S}_2\}$. We say that an interval-valued function $\mathcal{S}(\lambda)$ is *non-decreasing* (resp. *non-increasing*) in λ if $\mathcal{S}(\lambda) \preceq \mathcal{S}(\lambda')$ (resp $\mathcal{S}(\lambda') \preceq \mathcal{S}(\lambda)$) for $\lambda \leq \lambda'$.

In these terms, the interpolation of the soft constraints can be stated as follows:

Proposition B.2 (Soft constraints interpolate between **MaxUtil** and hard constrained solution). *Let $\Phi(t)$ be a convex, symmetric convex function with $\Phi(t) > 0$ for $t > 0$. Then the sets*

$$\begin{aligned} \mathcal{D}(\lambda) &:= \{\Delta := \langle \tau_A, \mathbf{w}_A \circ \pi_A \rangle - \langle \tau_B, \mathbf{w}_B \circ \pi_B \rangle : (\tau_A, \tau_B) \in \mathbf{P}(\lambda)\} \\ \mathcal{T}_A(\lambda) &:= \{t_A := \langle \tau_A, \mathbf{w}_A \circ \pi_A \rangle \mid \exists \tau_B : (\tau_A, \tau_B) \in \mathbf{P}(\lambda)\} \end{aligned}$$

are closed intervals. Moreover,

1. In all cases, $\lim_{\lambda \rightarrow \infty} \max\{|\Delta| \in \mathcal{D}(\lambda)\} = 0$.
2. If $0 \in \mathcal{D}(\lambda)$, then there exists a **MaxUtil** solution satisfying (18). Thus, for all $\lambda > 0$, $\mathbf{P}(\lambda) = \mathbf{P}(\infty)$.
3. If $\mathcal{D}(\lambda) \prec \{0\}$, then $\mathcal{D}(\lambda)$ and $\mathcal{T}_A(\lambda)$ are non-decreasing on $\lambda \in (0, \infty]$, and vice versa if $\mathcal{D}(\lambda) \succ \{0\}$.
4. If $\mathcal{D}(\lambda) \prec \{0\}$, then $\{0\} = \mathcal{D}(\infty) \succeq \mathcal{D}(\lambda) \succeq \{\min : \Delta \in \mathcal{D}(0)\}$, and $\mathcal{T}_A(\infty) \succeq \mathcal{T}_A(\lambda) \succeq \{\min : \Delta \in \mathcal{T}_A(\lambda)\}$, and vice versa if $\mathcal{D}(\lambda) \succ \{0\}$.

B.3.1 Proof of Proposition B.2

Again, we parameterize all solutions to the soft-constrained problem as in correspondence with solutions (t_A, t_B) to

$$\min_{(t_A, t_B)} g_A \mathcal{U}_A(t_A; \mathbf{w}_A) + g_B \mathcal{U}_B(t_B; \mathbf{w}_B) + \lambda \Phi(t_A - t_B).$$

Letting $\Delta := t_B - t_A$, we can reparameterize the above as

$$\min_{(t_A, \Delta)} g_A \mathcal{U}_A(t_A; \mathbf{w}_A) + g_B \mathcal{U}_B(t_A + \Delta; \mathbf{w}_B) - \lambda \Phi(\Delta).$$

Note then that $\mathcal{D}(\lambda)$ denotes the set of Δ which are partial maximizers of the above display. If $0 \in \{\mathcal{D}(\lambda)\}$, this implies that there exists a **MaxUtil** solution for which $\Delta = 0$, therefore, for all $\lambda > 0$, all solutions will be **MaxUtil** solutions for which $\mathcal{D}(\lambda) = 0$. Otherwise assume without loss of generality that $\mathcal{D}(\lambda) < \{0\}$.

First, the statement $\{0\} = \mathcal{D}(\infty) \succeq \mathcal{D}(\lambda) \succeq \{\min : \Delta \in \mathcal{D}(0)\}$, and $\mathcal{T}_A(\infty) \succeq \mathcal{T}_A(\lambda) \succeq \{\min : \Delta \in \mathcal{T}_A(\lambda)\}$, and vice versa if $\mathcal{D}(\lambda) \succ \{0\}$ can be solved by on a case-by-case basis. The strategy is to show that if any of these inequalities are violated, then the associated values of Δ and t_A are not partial maximizers of the soft constraint objective. In particular, $\mathcal{T}_A(\lambda) \subset [T_-, T_+]$ for some appropriate T_-, T_+ .

We now show that $\mathcal{D}(\lambda)$ and $\mathcal{T}_A(\lambda)$ are non-increasing and non-decreasing, respectively. We shall do so invoking the following technical lemma.

Lemma B.3. *Let $G_1(t)$ be concave and let $G_2(t; \lambda)$ be concave in t . Let $\partial G_2(t; \lambda)$ denote the super-gradient of G_2 , that is*

$$\partial G_2(t; \lambda) := \text{Conv}(\{\partial_- G_2(t; \lambda)\} \cup \{\partial_+ G_2(t; \lambda)\})$$

denotes the super-gradient set of the concave mapping $t \mapsto \partial G_2(t; \lambda)$.

Then if $\lambda \mapsto \partial G_2(t; \lambda)$ is non-increasing (resp. non-decreasing) in λ , the interval valued function defined below is non-increasing (resp. non-decreasing) in λ

$$\text{MAX}(\lambda) := \lambda \mapsto \arg \max_{t \in [a, b]} G_1(t) + G_2(t; \lambda).$$

For $\mathcal{D}(\lambda)$, one can write any partial maximizer Δ as

$$\max_{\Delta \geq 0} G_1(\Delta) + G_2(\Delta; \lambda)$$

with $G_1(\Delta) = \max_{t_A} g_A \mathcal{U}_A(t_A; \mathbf{w}_A) + g_B \mathcal{U}_B(t_A + \Delta; \mathbf{w}_B)$ and $G_2(\Delta; \lambda) = \lambda \Phi(\Delta)$. Note that $G_1(\Delta)$ is concave, being the partial maximization of a concave function, and $\partial G_2(\Delta; \lambda) = -t \partial \Phi(\Delta)$. Since $\partial \Phi(\Delta) \succeq \{0\}$ for $\Delta \geq 0$ (by convexity of ϕ), we have that $\partial G_2(\Delta; \lambda) = -t \partial \Phi(\Delta)$ is non-increasing in λ . Hence Lemma B.3 implies that interval valued function $\mathcal{D}(\lambda)$ is non-increasing.

To show that $\mathcal{T}_A(\lambda)$ is non-decreasing, we have that any maximizer t_A can be written as

$$\max_{t_A \in [T_-, T_+]} G_1(t_A) + G_2(t_A; \lambda)$$

where $G_1(t_A) = g_A \mathcal{U}_A(t_A; \mathbf{w}_A)$ and $G_2(t_A; \lambda) = \max_{\Delta \geq 0} g_B \mathcal{U}_B(t_A + \Delta; \mathbf{w}_B) + \lambda \Phi(\Delta)$. By Danskin's theorem,

$$\partial G_2(t_A; \lambda) = \{\partial \mathcal{U}_B(t_A + \Delta; \mathbf{w}_B) : \Delta \in \arg \max G_2(t_A; \lambda)\}.$$

Note that $\{\Delta \in \arg \max G_2(t_A; \lambda)\}$ is non-increasing in λ for a fixed t_A , since the contribution of the regularizer increases. Since the sets $\partial \mathcal{U}_B(t_A + \Delta; \mathbf{w}_B)$ are themselves non-increasing in Δ by

concavity, we conclude that $\partial G_2(t_A; \lambda)$ is non-decreasing in λ . Hence, Lemma B.3 implies that $\mathcal{T}_A(\lambda)$ is non-decreasing in λ .

Finally, to show that $\max\{|\Delta| : \Delta \in \mathcal{D}(\lambda)\} \rightarrow 0$, Note that the left and right derivatives of $g_A \mathcal{U}_A(t; \mathbf{w}_A)$ and $g_B \mathcal{U}_B(t; \mathbf{w}_B)$ are upper bounded by M whereas, since Φ is strictly convex, we know that for every $\epsilon > 0$, $\min\{|\partial_+ \Phi(\Delta)|, |\partial_- \Phi(\Delta)|\} > m(\epsilon)$ for all $\Delta : |\Delta| > \epsilon$. Hence, the first order optimality conditions cannot be satisfied for $|\Delta| > \epsilon$, and $\lambda > \frac{M}{m(\epsilon)}$, so as $\lambda \rightarrow \infty$, $|\Delta| \rightarrow 0$.

Proof of Lemma B.3. We prove the case where $\partial G_2(t; \lambda)$ is non-increasing. The first order conditions requires that at an optimal t , one has

$$\partial_- G_1(t) + \partial G_2(t; \lambda)_- \geq 0 \geq \partial_+ G_1(t) + \partial G_2(t; \lambda)_+$$

where the super-gradients are amended to take into account boundary conditions. Suppose that for the sake of contradiction that for $\lambda' > \lambda$, $\text{MAX}(\lambda') \preceq \text{MAX}(\lambda)$ fails. Then, there (a) exists a $t \in \text{MAX}(\lambda)$ such that $\{t\} \prec \text{MAX}(\lambda')$, or (b) $t \in \text{MAX}(\lambda')$ such that $\{t\} \succ \text{MAX}(\lambda)$. Note that if $\{t\} \prec \text{MAX}(\lambda')$, it must be the case that

$$\partial_+ G_1(t) + \partial G_2(t; \lambda')_+ > 0.$$

By assumption, $\partial_- G_2(t; \lambda')_+ \leq \partial_0 G_2(t; \lambda)_+$, which implies

$$\partial_+ G_1(t) + \partial G_2(t; \lambda')_+ \leq \partial_+ G_1(t) + \partial_+ G_2(t; \lambda)_- 0 \leq 0,$$

a contradiction. □

C Proofs of Main Results

We remark that the proofs in this section rely crucially on the characterizations of the optimal fairness-constrained policies developed in Section 6. We first define the notion of CDF domination, which is referred to in a few of the proofs. Intuitively, it means that for any score, the fraction of group B above this is higher than that for group A. It is realistic to assume this if we keep with our convention that group A is the disadvantaged group relative to group B.

Definition C.1 (CDF domination). π_A is said to be *dominated* by π_B if $\forall a \geq 1, \sum_{x>a} \pi_A < \sum_{x>a} \pi_B$. We denote this as $\pi_A \prec \pi_B$.

We remark that the $\prec$ notation in this section is entirely unrelated to the the partial order on intervals from Section B.3. Frequently, we shall use the following lemma:

Lemma C.1. *Suppose that $\pi_A \prec \pi_B$. Then, for all $\beta > 0$, it holds that $Q_A(\beta) \leq Q_B(\beta)$ and $\mathbf{u}(Q_A(\beta)) \leq \mathbf{u}(Q_B(\beta))$*

Proof. The fact that $Q_A(\beta) \leq Q_B(\beta)$ follows directly from the definition of monotonicity of $\mathbf{u}$ implies that $\mathbf{u}(Q_A(\beta)) \leq \mathbf{u}(Q_B(\beta))$. □

C.1 Proof of Proposition 3.1

The **MaxUtil** policy for group j solves the optimization

$$\max_{\tau_j \in [0,1]^C} \mathcal{U}_j(\tau_j) = \max_{\beta_j \in [0,1]} \mathcal{U}_j(r_{\pi_j}^{-1}(\beta_j)).$$

Computing left and right derivatives of this objective yields

$$\partial_+ \mathcal{U}_j(r_{\pi_j}^{-1}(\beta_j)) = \mathbf{u}(\mathbf{Q}_j(\beta)), \quad \partial_- \mathcal{U}_j(r_{\pi_j}^{-1}(\beta_j)) = \mathbf{u}(\mathbf{Q}_j^+(\beta)).$$

By concavity, solutions β^* satisfy

$$\begin{aligned} \beta < \beta^* & \quad \text{if} \quad \mathbf{u}(\mathbf{Q}_j(\beta)) > 0, \\ \beta > \beta^* & \quad \text{if} \quad \mathbf{u}(\mathbf{Q}_j^+(\beta)) < 0. \end{aligned} \tag{30}$$

Therefore, we conclude that the **MaxUtil** policy loans only to scores x s.t. $\mathbf{u}(x) > 0$, which implies $\Delta(x) > 0$ for all scores loaned to. Therefore we must have that $0 \leq \Delta \mu^{\text{MaxUtil}}$. By definition $\Delta \mu^{\text{MaxUtil}} \leq \Delta \mu^*$.

C.2 Proof of Corollary 3.2

We begin with proving part (a), which gives conditions under which **DemParity** cases relative improvement. Recall that $\bar{\beta}$ is the largest selection rate for which $\mathcal{U}(\bar{\beta}) = \mathcal{U}(\beta_A^{\text{MaxUtil}})$. First, we derive a condition which bounds the selection rate $\beta_A^{\text{DemParity}}$ from below. Fix an acceptance rate β such that $\beta_A^{\text{MaxUtil}} < \beta < \min\{\beta_B^{\text{MaxUtil}}, \bar{\beta}\}$. By Theorem 6.1, we have that **DemParity** selects to group A with rate higher than β as long as

$$g_A \leq g_1 := \frac{1}{1 - \frac{\mathbf{u}(\mathbf{Q}_A(\beta))}{\mathbf{u}(\mathbf{Q}_B(\beta))}}.$$

By (30) and the monotonicity of $\mathbf{u}$, $\mathbf{u}(\mathbf{Q}_A(\beta)) < 0$ and $\mathbf{u}(\mathbf{Q}_B(\beta)) > 0$, so $0 < g_1 < 1$.

Next, we derive a condition which bounds the selection rate $\beta_A^{\text{DemParity}}$ from above. First, consider the case that $\beta_B^{\text{MaxUtil}} < \bar{\beta}$, and fix β' such that $\beta_B^{\text{MaxUtil}} < \beta' < \bar{\beta}$. Then **DemParity** selects group A at a rate $\beta_A < \beta'$ for any proportion g_A . This follows from applying Theorem 6.1 since we have that $\mathbf{u}(\mathbf{Q}_A^+(\beta')) < 0$ and $\mathbf{u}(\mathbf{Q}_B^+(\beta')) < 0$ by (30) and the monotonicity of $\mathbf{u}$.

Instead, in the case that $\beta_B^{\text{MaxUtil}} > \bar{\beta}$, fix β' such that $\bar{\beta} < \beta' < \beta_B^{\text{MaxUtil}}$. Then **DemParity** selects group A at a rate less than β' as long as

$$g_A \geq g_0 := \frac{1}{1 - \frac{\mathbf{u}(\mathbf{Q}_A^+(\beta'))}{\mathbf{u}(\mathbf{Q}_B^+(\beta'))}}.$$

By (30) and the monotonicity of $\mathbf{u}$, $0 < g_0 < g_1$. Thus for $g_A \in [g_0, g_1]$, the **DemParity** selection rate for group A is bounded between β and β' , and thus **DemParity** results in relative improvement.

Next, we prove part (b), which gives conditions under which **EqOpt** cases relative improvement. First, we derive a condition which bounds the selection rate β_A^{EqOpt} from below. Fix an acceptance rate β such that $\beta_A^{\text{MaxUtil}} < \beta$ and $\beta_B^{\text{MaxUtil}} > G^{(A \rightarrow B)}(\beta)$. By Theorem 6.2, **EqOpt** selects group A

at a rate higher than β as long as

$$g_A > g_3 := \frac{1}{1 - \frac{1}{\kappa} \cdot \frac{\rho(Q_B(G^{(A \rightarrow B)}(\beta)))}{u(Q_B(G^{(A \rightarrow B)}(\beta)))} \frac{u(Q_A(\beta))}{\rho(Q_A(\beta))}}.$$

By (30) and the monotonicity of $\mathbf{u}$, $\mathbf{u}(Q_A(\beta)) < 0$ and $\mathbf{u}(Q_B(G^{(A \rightarrow B)}(\beta))) > 0$, so $g_3 > 0$.

Next, we derive a condition which bounds the selection rate β_A^{EqOpt} from above. First, consider the case that there exists β' such that $\beta' < \bar{\beta}$ and $\beta_B^{\text{MaxUtil}} < G^{(A \rightarrow B)}(\beta')$. Then EqOpt selects group A at a rate less than this β' for any g_A . This follows from Theorem 6.2 since we have that $\mathbf{u}(Q_A^+(\beta')) < 0$ and $\mathbf{u}(Q_B^+(G^{(A \rightarrow B)}(\beta')) < 0$ by (30) and the monotonicity of $\mathbf{u}$.

In the other case, fix β' such that $\beta < \beta' < \bar{\beta}$ and $\beta_B^{\text{MaxUtil}} > G^{(A \rightarrow B)}(\beta')$. By Theorem 6.2, EqOpt selects group A at a rate lower than β' as long as

$$g_A > g_2 := \frac{1}{1 - \frac{1}{\kappa} \cdot \frac{\rho(Q_B^+(G^{(A \rightarrow B)}(\beta')))}{u(Q_B^+(G^{(A \rightarrow B)}(\beta')))} \frac{u(Q_A^+(\beta'))}{\rho(Q_A^+(\beta'))}}.$$

By (30) and the monotonicity of $\mathbf{u}$, $0 < g_2 < g_3$. Thus for $g_A \in [g_2, g_3]$, the EqOpt selection rate for group A is bounded between β and β' , and thus EqOpt results in relative improvement.

C.3 Proof of Corollary 3.3

Recall our assumption that $\beta > \beta_A^{\text{MaxUtil}}$ and $\beta_B^{\text{MaxUtil}} > \beta$. As argued in the above proof of Corollary 3.2, by (30) and the monotonicity of $\mathbf{u}$, $\mathbf{u}(Q_A(\beta)) < 0$ and $\mathbf{u}(Q_B(\beta)) > 0$. Applying Theorem 6.1, DemParity selects at a higher rate than β for any population proportion $g_A \leq g_0$, where $g_0 = 1/(1 - \frac{1}{\kappa} \cdot \frac{\rho(Q_B(\beta))}{u(Q_B(\beta))}) \in (0, 1)$. In particular, if $\beta = \beta_0$, which we defined as the harm threshold (i.e. $\Delta\mu_A(r_{\pi_A}^{-1}(\beta_0)) = 0$ and $\Delta\mu_A$ is decreasing at β_0), then by the concavity of $\Delta\mu_A$, we have that $\Delta\mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{DemParity}})) < 0$, that is, DemParity causes active harm.

C.4 Proof of Corollary 3.4

By Theorem 6.2, EqOpt selects at a higher rate than β for any population proportion $g_A \leq g_0$, where $g_0 = 1/(1 - \frac{1}{\kappa} \cdot \frac{\rho(Q_B(G^{(A \rightarrow B)}(\beta)))}{u(Q_B(G^{(A \rightarrow B)}(\beta)))} \frac{u(Q_A(\beta))}{\rho(Q_A(\beta))})$. Using our assumptions $\beta_B^{\text{MaxUtil}} > G^{(A \rightarrow B)}(\beta)$ and $\beta > \beta_A^{\text{MaxUtil}}$, we have that $\mathbf{u}(Q_B(G^{(A \rightarrow B)}(\beta))) > 0$ and $\mathbf{u}(Q_A(\beta)) < 0$, by (30) and the monotonicity of $\mathbf{u}$. This verifies that $g_0 \in (0, 1)$. In particular, if $\beta = \beta_0$, then by the concavity of $\Delta\mu_A$, we have that $\Delta\mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{EqOpt}})) < 0$, that is, EqOpt causes active harm.

C.5 Proof of Corollary 3.5

Applying Theorem 6.1, we have

$$-\frac{1 - g_A}{g_A} \mathbf{u}(Q_A(\beta)) < \mathbf{u}(Q_B(\beta)) \implies \beta_{\text{DemParity}} > \beta.$$

Applying Theorem 6.2, we have:

$$\mathbf{u}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta))) \cdot \frac{\langle \boldsymbol{\rho}, \boldsymbol{\pi}_B \rangle}{\langle \boldsymbol{\rho}, \boldsymbol{\pi}_A \rangle} \cdot \frac{\boldsymbol{\rho}(\mathbf{Q}_A^+(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_B^+(G^{(A \rightarrow B)}(\beta)))} < -\frac{1 - g_A}{g_A} \mathbf{u}(\mathbf{Q}_A^+(\beta)) \implies \beta_{\text{EqOpt}} < \beta.$$

By Corollaries 3.3 and 3.4, choosing $g_A < g_2 := 1/(1 - \frac{\mathbf{u}(\mathbf{Q}_A(\beta))}{\mathbf{u}(\mathbf{Q}_B(\beta))})$ and $g_A > g_1 := 1/(1 - \frac{1}{\kappa} \cdot \frac{\boldsymbol{\rho}(\mathbf{Q}_B^+(G^{(A \rightarrow B)}(\beta)))}{\mathbf{u}(\mathbf{Q}_B^+(G^{(A \rightarrow B)}(\beta)))} \cdot \frac{\mathbf{u}(\mathbf{Q}_A^+(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_A^+(\beta))})$ satisfies the above.

It remains to check that $g_1 < g_2$. Since we assumed $\beta > \sum_{x > \mu_A} \pi_A$, we may apply Lemma C.2 to verify this.

Thus we indeed have sufficient conditions for $\beta_{\text{DemParity}} > \beta > \beta_{\text{EqOpt}}$. In particular, if $\beta = \beta_0$, then by the concavity of $\Delta \mu_A$, we have that $\Delta \mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{EqOpt}})) > 0$, that is, EqOpt causes improvement, and $\Delta \mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{DemParity}})) < 0$, that is, DemParity causes active harm.

Lastly, because $\beta_{\text{DemParity}} > \beta_{\text{EqOpt}}$, it is always true that $\Delta \mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{DemParity}})) > 0 \implies \Delta \mu_A(r_{\pi_A}^{-1}(\beta_A^{\text{EqOpt}})) > 0$, using the concavity of the outcome curve.

Lemma C.2 (Comparison of DemParity and EqOpt selection rates). *Fix $\beta \in [0, 1]$. Suppose π_A, π_B are identical up to a translation with $\mu_A < \mu_B$. Also assume $\boldsymbol{\rho}(x)$ is affine in x . Denote $\kappa = \frac{\langle \boldsymbol{\rho}, \boldsymbol{\pi}_B \rangle}{\langle \boldsymbol{\rho}, \boldsymbol{\pi}_A \rangle}$. Then,*

$$\beta > \sum_{x > \mu_A} \pi_A$$

$$\text{implies } \mathbf{u}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta))) \cdot \kappa \cdot \frac{\boldsymbol{\rho}(\mathbf{Q}_A(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta)))} < \mathbf{u}(\mathbf{Q}_B(\beta)).$$

Proof. If we have $\beta > \sum_{x > \mu_A} \pi_A$, by lemma C.3, we must also have $\frac{\mu_B}{\mu_A} < \frac{\mathbf{Q}_B(\beta_0)}{\mathbf{Q}_A(\beta_0)}$. This implies $\kappa = \frac{\sum_x \pi_B(x) \boldsymbol{\rho}(x)}{\sum_x \pi_A(x) \boldsymbol{\rho}(x)} < \frac{\boldsymbol{\rho}(\mathbf{Q}_B(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_A(\beta_0))}$ by linearity of expectation and linearity of $\boldsymbol{\rho}$. Therefore,

$$\kappa \cdot \frac{\boldsymbol{\rho}(\mathbf{Q}_A(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_B(\beta_0))} < 1 \quad (31)$$

Further, using $G^{(A \rightarrow B)}(\beta) > \beta$ from lemma C.3 and the fact that $\frac{\mathbf{u}(x)}{\boldsymbol{\rho}(x)}$ is increasing in x , we have $\frac{\mathbf{u}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta)))}{\boldsymbol{\rho}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta)))} < \frac{\mathbf{u}(\mathbf{Q}_B(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_B(\beta))}$. Therefore, $\mathbf{u}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta))) \cdot \kappa \cdot \frac{\boldsymbol{\rho}(\mathbf{Q}_A(\beta_0))}{\boldsymbol{\rho}(\mathbf{Q}_B(G^{(A \rightarrow B)}(\beta_0)))} < \kappa \cdot \frac{\mathbf{u}(\mathbf{Q}_B(\beta))}{\boldsymbol{\rho}(\mathbf{Q}_B(\beta))} \cdot \boldsymbol{\rho}(\mathbf{Q}_A(\beta)) < \mathbf{u}(\mathbf{Q}_B(\beta))$ where the last inequality follows from (31). $\square$

We use the following technical lemma in the proof of the above lemma.

Lemma C.3. *If π_A, π_B that are identical up to a translation with $\mu_A < \mu_B$, then*

$$G(\beta) > \beta \quad \forall \beta, \quad (32)$$

$$\beta > \sum_{x > \mu} \pi_A \implies \frac{\mu_B}{\mu_A} < \frac{\mathbf{Q}_B(\beta)}{\mathbf{Q}_A(\beta)}. \quad (33)$$

Proof. For (32), observe that $\text{TPR}_A = \boldsymbol{\rho}(\mu_A) < \text{TPR}_B = \boldsymbol{\rho}(\mu_B)$. For any β , we can write $\mathbf{Q}_B(\beta) = \mu_B + c$ and $\mathbf{Q}_A(\beta) = \mu_A + c$ for some c , since π_A, π_B that are identical up to translation by

$\mu_A - \mu_B$. Thus, by computation, we can see that for $Q(\beta) < \mu$, $\partial_+ G^{(A \rightarrow B)}(\beta) > 1$ and for $Q(\beta) < \mu$, $\partial_+ G^{(A \rightarrow B)}(\beta) < 1$. Since $G^{(A \rightarrow B)}$ is monotonically increasing on $[0, 1]$, we must have $G^{(A \rightarrow B)}(\beta) > \beta$ for every $\beta \in [0, 1]$.

For (33), we have $\beta > \sum_{x > \mu} \pi_A$, we can again write $Q_B(\beta) = \mu_B - c$ and $Q_A(\beta) = \mu_A - c$, for some $c > 0$. Then it is clear than we have $\frac{\mu_B}{\mu_A} < \frac{Q_B(\beta)}{Q_A(\beta)}$. $\square$

C.6 Proof of Corollary 3.6

Proof. $\beta_A^{\text{MaxUtil}} < \beta_B^{\text{MaxUtil}}$ implies $g_A \cdot u(Q_A(\beta_A^{\text{MaxUtil}})) + g_B \cdot u(Q_B(\beta_A^{\text{MaxUtil}})) > 0$, which by Theorem 6.1, implies $\beta_A^{\text{MaxUtil}} < \beta_A^{\text{DemParity}}$.

$\text{TPR}_A(\tau^{\text{MaxUtil}}) > \text{TPR}_B(\tau^{\text{MaxUtil}})$ implies $G^{(A \rightarrow B)}(\beta_A^{\text{MaxUtil}}) > \beta_B^{\text{MaxUtil}}$ and so

$u(Q_B(G^{(A \rightarrow B)}(\beta_A^{\text{MaxUtil}}))) < 0$. Therefore by Theorem 6.2, we have that $\beta_A^{\text{MaxUtil}} > \beta_A^{\text{EqOpt}}$. $\square$

We now give a very simple example of $\pi_A \prec \pi_B$ where Theorem 3.5 holds. The construction of the example exemplifies the more general idea of using large in-group inequality in group A to skew the true positive rate at MaxUtil, making $\text{TPR}_A(\tau^{\text{MaxUtil}}) > \text{TPR}_B(\tau^{\text{MaxUtil}})$.

Example C.1 (EqOpt causes relative harm). Let $C = 6$, and let the utility function be such that $u(4) = 0$. Suppose $\pi_A(5) = 1 - 2\epsilon$, $\pi_A(1) = 2\epsilon$ and $\pi_B(5) = 1 - \epsilon$, $\pi_B(3) = \epsilon$.

We can easily check that $\pi_A \prec \pi_B$. However, for any $\epsilon \in (0, 1/4)$, we have that $\text{TPR}_B(\tau^{\text{MaxUtil}}) = \frac{5(1-\epsilon)}{5(1-\epsilon)+3\epsilon} < \text{TPR}_A(\tau^{\text{MaxUtil}}) = \frac{5(1-2\epsilon)}{5(1-2\epsilon)+2\epsilon}$.

C.7 Proof of Proposition 4.1

Denote the upper quantile function under $\hat{\pi}$ as $\hat{Q}$. Since $\hat{\pi} \prec \pi$, we have $\hat{Q}(\beta) \leq Q(\beta)$. The conclusion follows for MaxUtil and DemParity from Theorem 6.1 by the monotonicity of u .

If we have that $\text{TPR}_A(\tau) > \widehat{\text{TPR}}_A(\tau) \forall \tau$, that is, the true TPR dominates estimated TPR, the conclusion for EqOpt follows from Theorem 6.2, by the same argument as in the proof of Corollary 3.6.

C.8 Proof of Proposition 4.2

By Proposition 5.3, $\beta^* = \arg\max_{\beta} \Delta\mu_A(\beta)$ exists and is unique. $\beta_0 = \max\{\beta \in [\beta_A^{\text{MaxUtil}}, 1] : \mathcal{U}(\beta_A^{\text{MaxUtil}}) - \mathcal{U}_A(\beta) \leq \delta\}$ which exists and is unique, by the continuity of $\Delta\mu_A$ and Proposition 5.3.