

---

# Strategic ranking

---

Lydia T. Liu

University of California, Berkeley

Nikhil Garg

Cornell Tech and the Technion

Christian Borgs

University of California, Berkeley

## Abstract

Strategic classification studies the design of a classifier robust to the manipulation of input by strategic individuals. However, the existing literature does not consider the effect of *competition* among individuals as induced by the algorithm design. Motivated by constrained allocation settings such as college admissions, we introduce *strategic ranking*, in which the (designed) individual reward depends on an applicant’s post-effort rank in a measurement of interest. Our results illustrate how competition among applicants affects the resulting equilibria and model insights. We analyze how various ranking reward designs, belonging to a family of step functions, trade off applicant, school, and societal utility, as well as how ranking design counters inequities arising from disparate access to resources. In particular, we find that randomization in the reward design can mitigate two measures of disparate impact, welfare gap and access.

## 1 INTRODUCTION

Many allocative decisions—from education to employment—rely on relative quality across individuals, not absolute quality: Berkeley accepts  $\approx 15\%$  of college applicants per year, major CS conferences accept about  $\approx 20\%$  of submissions, and each job opening is filled by one candidate. Even if every applicant improves substantially, due to (perceived) capacity constraints the number accepted would not commensurately increase. The comparative aspect of *ranking* differs sharply from *classification*, which evaluates each entry in isolation. While this difference is often overlooked in the ML community, we

illustrate that it is consequential and motivates *access* as a measure of disparate impact from algorithmic decisions.

Our specific context is the literature on *strategic classification*, which aims to address gaming by applicants controlling the classifier’s inputs (Brückner et al., 2012; Hardt et al., 2016)—an instance of Goodhart’s law. For example, if a predictive classifier for admissions learns that a student’s number of extracurricular activities correlates with college achievement and thus bases admissions decisions on it, students may list many more activities on their applications without devoting any time to them. A strategic classifier aims to undo this effect, informally speaking, by shifting weight towards features that are costly to game, such that the designer’s utility is maximized even taking into account strategic behavior—as in a Stackelberg equilibrium.

Importantly, many motivating applications in the strategic classification literature, such as manipulating test scores to gain college admission, are ranking problems (or classification with capacity constraints on how many can be classified with each label). This aspect induces competition between applicants, affecting their effort and in turn their comparative performance and ultimate relative position in the ranking. To study the interaction of these effects, this work introduces the problem of *strategic ranking*: we study the general welfare effects of reward design for ranking, and examine the role of randomization in trading off designer’s utility and population welfare, as well as reducing disparate impact. We study such *competition* between applications by shifting the task from classifier design to *ranking reward design*, and so fill a gap in the current literature on algorithmic fairness and strategic behavior that has thus far focused on classification (and regression).

In pursuit of these questions, our theoretical framework recalls the long-standing economics literature on *contests*, in which agents expend effort to obtain a reward that is a function of their relative performance across agents (Bodoh-Creed and Hickman, 2018; Olszewski and Siegel, 2016, 2019a,b), as well as that on

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

signaling and strategic behavior (Spence, 1978). We elaborate on the connections and distinctions with this literature in Section 1.1.

In our strategic ranking framework, a designer ranks applicants by a single, observed measure, called a score. The score is a function of the applicant’s (possibly multi-dimensional) chosen effort level(s), latent skill(s), and environment. The designer controls the reward  $\lambda(\theta)$  assigned to each rank  $\theta$ , under an overall reward constraint; e.g., in our primary interpretation, the reward is a probability at which an applicant of each rank is admitted. Realistically, there may be complexity constraints on the ranking reward function  $\lambda$  that the designer can deploy—in this work, we analyze settings where  $\lambda$  belongs to a family of step functions that we call “K-level policies”. In response, the applicants choose their (costly) effort level for the observed measure, such that in equilibrium—that is, given the effort levels of the entire applicant population—their effort level and resulting rank maximizes their welfare (reward minus effort cost). Depending on the setting, the designer’s objective may be a function of applicant effort and who is admitted. We assume that applicant effort improves the designer’s utility function, but not their own welfare except indirectly through the ranking reward; in other words, applicants are not intrinsically motivated to exert effort. Using this model, our contributions are:

**First** (Section 2), we analyze the equilibrium behavior and resulting optimal designs, illuminating important differences between the classification and ranking settings. We show that, in a general setting and with any reasonable reward function  $\lambda$ , the competition effect results in *rank preservation*: in equilibrium, the ranks (and thus rewards) of applicants after applicant effort is the same as before. This finding differs from strategic classification – where one must adapt to the effort it induces to maintain accuracy (Hardt et al., 2016) – and simplifies equilibrium analysis.

**Second** (Section 3), we study how the design of reward function  $\lambda$  differentially affects the welfare of various stakeholders: applicants, a school preferring to admit those with the highest score, and a social planner maximizing the score over the population. We find, e.g., there is an trade-off even in one skill dimension: while deterministically admitting the highest ranked students maximizes the school’s utility (among *two*-level policies but not generally), it leads to applicants exerting costly effort. Adding randomness reduces pressure for applicants at the expense of the school’s utility.

**Third** (Section 4) we analyze equilibria in the presence of structural inequities between different groups, finding that competition amplifies such inequities without

careful reward design: designing reward  $\lambda$  to increasing school utility increases the welfare gap between groups, and decreases *access* (group specific admission probability) for the disadvantaged group.

Our definitions of school utility and group access—natural in constrained allocation settings—have not been studied in unconstrained settings, where policies may admit different numbers of applicants. We also extend our model to a multi-dimensional setting in Appendix B.

### 1.1 Related work

Our work sits at the intersection of work in two communities, *strategic classification* and *fair machine learning* in computer science, and *contests* and *effort* in economics.

**Strategic classification** As in strategic classification (Brückner et al., 2012; Dong et al., 2018; Hardt et al., 2016), we consider the challenge an institution faces when deploying a classifier that applicants can game; a naive classifier that does not factor in the resulting distribution shifts would be inaccurate. Unlike strategic classification, which considers the manipulation of observable features to be unproductive from the institution’s perspective, our model of efforts stipulates that the school prefers higher effort levels as they give rise to higher scores. We follow the line of work considering the design of reward functions that further incentivize agent effort on productive tasks (Bechavod et al., 2021; Kleinberg and Raghavan, 2019; Miller et al., 2020; Shavit et al., 2020). We in particular draw inspiration from and compare to several recent directions in the strategic classification literature: with multiple agents (Alon et al., 2020; Haghtalab et al., 2020), when a classifier may be random (Braverman and Garg, 2020), and with fairness concerns (Hu et al., 2019; Milli et al., 2019). Like other strategic classification work, our work is further related to the prior economics literature on contract design, elaborated in Appendix B.

**Economics: Contests and modeling effort** Our theoretical analysis is similar to that in the literature on *contests* in economics (Barut and Kovenock, 1998; Bodoh-Creed and Hickman, 2018; Connelly et al., 2014; Fang et al., 2020; Konrad, 2007; Olszewski and Siegel, 2016, 2019a,b); e.g., our rank preservation result in Section 2.2 reflects assortative allocation results in the literature under similar conditions. The field is too extensive to summarize here, so we refer the reader to surveys (Corchón, 2007; Fu and Wu, 2019) and discuss the works closest to ours. Bodoh-Creed and Hickman (2018) develop a model with students endogenously choosing effort; they use the model to

compare various affirmative action schemes. Most related is work by Olszewski and Siegel (2016, 2019a,b). Olszewski and Siegel (2019b) find that policies that “pool” individuals into tiers (as opposed to continuous ranks) can be Pareto improving for students; even those with decreased individual rewards would benefit from decreased competitive pressure.

Our work departs from the contests literature through its focus on the questions most common in the strategic classification: while the former primarily considers the efforts and resulting welfare of participants, we study how the design of the reward function  $\lambda$  differentially affects the designer’s utility, applicant welfare, and fairness metrics. In particular, when  $\lambda$  is interpreted as a probability of admission, our designer faces trade-offs with incentivizing overall effort and admitting the most skilled students (Section 3); with disparate access, the designer must further navigate trade-offs between fairness and inducing effort (Section 4); with multiple score dimensions, the designer must ensure that applicants do not excessively game one dimension at the cost of the other (Appendix B). Simultaneous to us, Elkind et al. (2021) consider the effect of contest design on participant and designer welfare.

Beyond contests, our work connects to the economics literature on effort and subsequent reward (Becker, 1973; Calsamiglia, 2009; Roemer, 1998; Spence, 1978).

**FATE in machine learning and MD4SG** We broadly connect to the Fairness, Accessibility/Accountability, Transparency, and Ethics in machine learning (Chouldechova and Roth, 2018) and the Mechanism Design for Social Good literatures (Abebe and Goldner, 2018). Most relevant are works on fair ranking (Mathioudakis et al., 2020; Tabibian et al., 2020; Zehlike and Castillo, 2020; Zehlike et al., 2017), constrained allocation (Aziz, 2020; Cai et al., 2020; Gözl et al., 2019; Noriega-Campero et al., 2019), and admissions (Faenza et al., 2020; Garg et al., 2020; Hu and Chen, 2018; Immorlica et al., 2019; Kannan et al., 2019, 2021; Liu et al., 2018, 2020; Liu and Garg, 2021; Mouzannar et al., 2019; Rolf et al., 2020). We consider such effects as they interact with agents’ strategic responses to the mechanism.

## 2 MODEL AND EQUILIBRIUM

Each *applicant* has an (unobserved) skill level and so a pre-effort *rank*. The applicants are modeled as a continuum of players (Schmeidler, 1973), rather than a finite number of atoms; informally, applicants do not respond to the strategies of every other applicant combinatorially, but rather the applicant population as a

whole. Applicants choose effort levels, resulting in post-effort scores and ranks. A single *school* determines rewards for each post-effort rank level, thus affecting applicant incentives to choose their effort.

### 2.1 Model

**Applicants** There is a unit mass of *applicants*, indexed by an observed index  $\omega \in [0, 1]$  distributed uniformly.<sup>1</sup> Each applicant has a latent (unobserved) skill level represented by some measurable function of  $\omega$ . We assume that the distribution of the skills has no atoms and that the CDF of this distribution is strictly increasing. Using the CDF to map the skill of an applicant to a rank in  $[0, 1]$ , each applicant gets an (unobserved) rank  $\theta_{\text{pre}} = \theta_{\text{pre}}(\omega)$  which by our assumption on the CDF is again uniformly distributed in  $[0, 1]$  (the higher the rank the better). With this setup, the skill of an applicant with rank  $\theta_{\text{pre}}$  can be written as  $f(\theta_{\text{pre}})$  where  $f$  is a strictly increasing, continuous function. It will be notationally convenient to label applicants by their rank  $\theta_{\text{pre}}$ , though the reader should note that in contrast to  $\omega$ ,  $\theta_{\text{pre}}(\omega)$  is assumed to be unobservable.

Each applicant chooses an *effort* level  $e \geq 0$ , the result of which is an observed, post-effort *score*,  $v = v(e, \theta_{\text{pre}}) = g(e) \cdot f(\theta_{\text{pre}})$ . In other words, we assume the post-effort score is the product of two components, associated with the effort and the pre-effort skill respectively. We assume that the effort transfer function  $g : [0, \infty) \mapsto [0, \infty)$  is a continuous, concave, strictly increasing function, representing that marginal effort improves one’s score but has diminishing returns. The strategies of the applicants can then be described by a function  $\theta_{\text{pre}} \mapsto e(\theta_{\text{pre}})$ . Each applicant is then ranked according to their score  $v$ , resulting in a *post-effort rank*  $\theta_{\text{post}}$ . Note that the ranking  $\theta_{\text{post}}$  is slightly less trivial than a ranking of the skills, since the scores might have ties, which have to be resolved.

**Tie Breaking.** Given a choice of strategies  $\theta_{\text{pre}} \mapsto e(\theta_{\text{pre}})$ , let  $F$  be the CDF of the scores  $v(e(\theta_{\text{pre}}), \theta_{\text{pre}})$ . Since atoms for the distribution of  $v$  would lead to ties for ranks defined as  $F(v)$ , we will use the labels of the applicants to break ties with the help of a (publicly announced) tie-breaking function  $\Gamma(\omega)$ , defined, e.g., via a collision free hash of the applicant names. Here we require that  $\Gamma$  is a measurable function from  $[0, 1]$  to  $[0, 1]$  that maps different applicant labels to different values. We then use  $\Gamma$  to resolve the atoms of  $F$ , leading to a ranking function  $v \mapsto \gamma(\omega, v)$  which is equal to  $F(v)$  except when  $v$  is an atom of the score distribution, in which case it takes values in the “gap

<sup>1</sup>The index  $\omega$  should be interpreted as each applicant’s “name,” uncorrelated with skill, used solely for tie-breaking.

interval”  $[F_-(v), F(v)]$ , where  $F_-(v)$  is the left limit of  $F$  at  $v$ . We construct  $\gamma$  in such a way that it gives the uniform distribution for  $\theta_{\text{post}}(\omega)$  if we set  $\theta_{\text{post}}(\omega) = \gamma(\omega, v(e(\theta_{\text{pre}}(\omega)), \theta_{\text{pre}}(\omega)))$ .<sup>2</sup>

**Ranking designer (school)** A single *school* is admitting applicants, based on their ranking. In particular, the school can choose a ranking reward function  $\lambda : [0, 1] \mapsto [0, 1]$ , such that an applicant with post-effort rank  $\theta_{\text{post}}$  is admitted with probability  $\lambda(\theta_{\text{post}})$ . We assume that  $\lambda$  is non-decreasing and that the school has a constraint on the overall probability, such that in expectation it admits a number of applicants equal to a capacity constraint  $\rho \in (0, 1)$ , i.e.,  $\mathbb{E}_{\theta_{\text{post}}}[\lambda(\theta_{\text{post}})] = \rho$ .<sup>3</sup> We may also refer to  $\lambda$ , informally, as the admission policy.

For simplicity, we further assume that  $\lambda$  is a step-function with  $K$  distinct levels  $\ell_0 < \dots < \ell_{K-1}$ , and  $K - 1$  cut-points parameterized by  $c_1 < \dots < c_{K-1}$  (with  $c_0 = 0, c_K = 1$ ). In other words, we have  $\lambda(\theta) = \ell_k$ , for all  $\theta \in \psi_k \triangleq [c_k, c_{k+1})$ . Thus, applicants in the same post-rank interval  $\theta \in \psi_k$  receive the same reward.

### Individual applicant welfare and equilibrium

Given the designer’s function  $\lambda$  and the effort levels of other applicants, each applicant chooses effort  $e$  to maximize their individual welfare,

$$W(e, \lambda(\theta_{\text{post}})) = \lambda(\theta_{\text{post}}) - p(e),$$

where the effort cost function  $p$  is non-negative, continuous, and strictly convex on  $[0, \infty)$ , with  $e_0 := \arg\min_e p(e)$  and  $p(e_0) = 0$ .

Applicants are assumed not to personally benefit from increasing their score  $v$ , except through the corresponding increase in their rank and reward. While the definition of  $p$  does not preclude the applicant receiving any intrinsic benefit from exerting effort, we assume that the *net* benefit from effort is non-positive.

After a school chooses its ranking reward function  $\lambda$ , each applicant chooses their effort level. However, unlike in strategic classification, in the ranking setting applicants must further take into account the effort levels (and resulting post-effort values) of other applicants. An *equilibrium* of effort levels is then an assignment  $\theta_{\text{pre}} \mapsto e(\theta_{\text{pre}})$  of effort levels and resulting post-effort rank rewards in which given the efforts of other applicants, no applicant can increase their welfare by changing their effort. This is formalized in the following definition.

<sup>2</sup>See Remark A.1 for details on this construction.

<sup>3</sup>We use  $\mathbb{E}_{\theta_{\text{post}}}[\cdot]$  to denote an integral over  $\theta_{\text{post}}$  (and  $\mathbb{E}[\cdot]$  to denote an integral over  $\omega$ ) with respect to the Lebesgue measure; informally, this can be thought of as averaging over the applicant population.

**Definition 1 (Equilibrium).** *Given a tie-breaking function  $\Gamma(\omega)$  and a ranking probability function  $\lambda$ , an equilibrium is a set of effort levels and post-effort ranking rewards for each applicant,  $\{e(\theta_{\text{pre}}), \lambda(\theta_{\text{post}}(\theta_{\text{pre}}))\}$  such that, for all  $\omega$ ,*

$$e(\theta_{\text{pre}}(\omega)) \in \arg\max_e W(e, \lambda(\gamma(\omega, v(e, \theta_{\text{pre}}(\omega)))))$$

$$\theta_{\text{post}}(\theta_{\text{pre}}(\omega)) = \gamma(\omega, v(e(\theta_{\text{pre}}(\omega)), \theta_{\text{pre}}(\omega))),$$

where  $\gamma(\omega, v)$  is the ranking induced by the CDF of the scores resulting from effort levels  $\{e(\theta_{\text{pre}})\}$  and the tie-breaking function  $\Gamma$ .<sup>4</sup>

In equilibrium, the strategy of the applicants is thus characterized by a collection of efforts  $e(\theta_{\text{pre}})$  and post-effort ranking rewards  $\lambda(\theta_{\text{post}}(\theta_{\text{pre}}))$  of the applicants as a function of  $\theta_{\text{pre}}$ , with a corresponding joint distribution for  $e$  and  $\theta_{\text{post}}$  induced by the equilibrium and the underlying uniform distribution over  $\theta_{\text{pre}}$ .<sup>5</sup> Intuitively, we need to define index  $\omega$  and the tie-breaking ordering  $\Gamma$  because while there may be ties in post-effort values, post-effort *ranks* must be unique. However, for notational ease, in the rest of this work we drop the index  $\omega$  and instead refer to applicants by their pre-effort rank  $\theta_{\text{pre}}$ . Our results hold for any ordering  $\Gamma$ , and so we further omit it.

**Aggregate welfare and utility** We define three aggregate welfare functions of the equilibrium efforts and scores to capture the interests of different stakeholders. The **applicant welfare** is defined as the population average of the individual applicant welfare at equilibrium:

$$\mathcal{W} := \mathbb{E}[W(e, \lambda(\theta_{\text{post}}))] = \rho - \mathbb{E}[p(e)].$$

On the other hand, society derives value from the scores of applicants post-effort, leading to the following **societal utility**:  $\mathcal{U}^{\text{soc}} := \mathbb{E}[v]$ .

In other words, society prefers the entire applicant population to achieve higher scores, not only those who are admitted to the school, since higher test scores are correlated with labor productivity and economic growth (Hanushek and Woessmann, 2010).

Finally, the design  $\lambda$  is controlled by a school, who may only draw value from those who enroll. The school’s **private utility** is the expected score of admitted applicants, which in our continuum formulation is the expectation of  $v$  weighted by  $\lambda(v)$ :  $\mathcal{U}^{\text{pri}} := \mathbb{E}[v \cdot \lambda(\theta_{\text{post}})]$ .

**Discussion** Our three welfare functions represent the utilities of three stakeholders in any ranking setting: the applicants, ranking institution, and broader society. As we show in Section 3, optimal design differs

<sup>4</sup>See Remark A.3 on  $\gamma$  and the set of efforts.

<sup>5</sup>See Remark A.4 on credible commitment by the school.

substantially for the three – the school maximizing its own utility comes at a cost to the others. Furthermore, we note that this comparison is possible *because* our setting is one of constrained allocation, in which the expected number of admitted applicants (integral over  $\lambda$ ) is fixed. For example, defining school’s private utility as the expected score of admitted applicants (and applicant welfare as containing admissions probability) is unnatural without such a constraint, as various mechanisms may admit (classify as ‘1’) different numbers of applicants; for this reason, strategic classification papers often consider classification accuracy as the school’s objective, which has no analogue in practice.

Our base model is purposely minimalist, to emphasize the competitive aspect of ranking and constrained allocation. We consider extensions to disparate access to resources between socioeconomic groups (Section 4) and multi-dimensional, potentially unobserved scores (Appendix B) after studying welfare under our base model.

## 2.2 Equilibria characterization

The definition of equilibria suggests that studying their properties may be difficult: our utility functions depend on both the efforts of applicants and their induced relative rankings, the relationship of which may be complex in general. However, the following result, reminiscent of much of the contests literature under similar assumptions (see Remark A.6), establishes that ranking rewards are preserved under effort.

**Proposition 2.1** (Rank preservation). *In every equilibrium,  $\lambda(\theta_{post}(\theta_{pre})) = \lambda(\theta_{pre})$ , up to sets of measure 0.*

As shown in the next section, rank preservation simplifies substantially the evaluation of various utility functions, as the post-effort rankings of applicants are fixed and known.

The result follows from the shared cost function  $p$  being convex and increasing on  $[e_0, \infty)$ : if an applicant pre-effort rank  $\theta$  finds it optimal to achieve post-effort value  $v$ , then each applicant with pre-effort rank  $\bar{\theta} > \theta$  finds it optimal to reach a post-effort value  $\bar{v} \geq v$ . We note that this result (and all our subsequent results) hold for any tie-breaking function  $\Gamma$ , as any set of post-effort values in which tie-breaking occurs across two or more reward bands  $\psi$  cannot constitute an equilibrium.

While the ranking reward function  $\lambda$  does not affect the induced rankings of applicants, it does determine the *effort* exerted by each, as formalized next.

**Theorem 2.2** (Second price effort). *There exists a*

Equilibrium efforts and scores (4-level policy)




Figure 1: Equilibrium efforts and scores for a sample 4-level ranking reward function.

*equilibrium such that  $\lambda(\theta_{post}(\theta_{pre})) = \lambda(\theta_{pre})$  and applicants with  $\theta_{pre} \in \psi_k \triangleq [c_k, c_{k+1})$  exert effort  $e_k(\theta_{pre})$ , where*

$$e_k(\theta_{pre}) = \begin{cases} e_0 & \text{for } k = 0 \\ \max \left( g^{-1} \left( \frac{g(\tilde{e}_{k-1}) \cdot f(c_k)}{f(\theta_{pre})} \right), e_0 \right) & \text{otherwise,} \end{cases}$$

*with  $\tilde{e}_{k-1} \geq e_{k-1}(c_k)$  inductively defined by*

$$p(\tilde{e}_{k-1}) = p(e_{k-1}(c_k)) + \ell_k - \ell_{k-1}$$

*The equilibrium is unique up to sets of measure 0.*

As the theorem name suggests, the effort exerted by each applicant in equilibrium is akin to the price paid in a second price auction (cf. Myerson (1981)): each applicant exerts just enough effort that applicants in the level below (those with ranks at or below the given applicant’s level’s lower cut-point  $c_k$ ) cannot increase their welfare by instead exerting additional effort,  $\tilde{e}_{k-1} > e_0$ . It is instructive to calculate the post-effort scores resulting from these efforts; for  $\theta_{pre} \in \psi_k$ , they are

$$v_k(\theta_{pre}) = \begin{cases} g(e_0)f(\theta_{pre}) & \text{for } k = 0 \\ \max \left\{ g(\tilde{e}_{k-1})f(c_k), g(e_0)f(\theta_{pre}) \right\} & \text{o.w.,} \end{cases}$$

Thus at the beginning of each band, applicants exert strictly more effort than those at the top of the previous band, decreasing their effort within the band since their increased skill requires less effort to get the same score. If at some point effort  $e_0$  (with cost  $p(e_0) = 0$ ) is enough to maintain the score needed to get reward  $\ell_k$ , their effort stays constant and their score grows, up to the beginning of the next band, when at least initially, the scores again stay constant and the efforts

decrease. Figure 1 illustrates the equilibrium efforts and post-effort scores under a sample 4-level reward function, and choices of  $f, g$  and  $p$  such that it is never enough to just exert effort  $e_0$  except in the first band.<sup>6</sup>

Thus, while the reward function  $\lambda$  does not change the ranking of applicants (by Lemma 2.1), it does affect the effort exerted and thus their post-effort scores. The function  $\lambda$  is thus a design parameter for the school whose objective depends on applicant scores, as we explore in the next section.<sup>7</sup>

Remark A.6 details the technical differences between our results and analogues in the contests literature (Bodoh-Creed and Hickman, 2018; Olszewski and Siegel, 2016).

### 3 WELFARE ANALYSIS

We begin our analysis by studying how the design of ranking reward function  $\lambda$  changes the utilities of the applicants, the school, and the society.

Recall that the school is constrained to a  $K$ -level  $\lambda$  admission policy such that the average probability of admission is  $\rho > 0$ , that is,  $\mathbb{E}[\lambda(\theta_{\text{post}})] = \rho$ . A special case of interest is the following *two-level* ( $K = 2$ ) function.

**Definition 2** (Two-level policy). *In our baseline two-level function parameterized by cut-off  $c \in (0, 1 - \rho]$ , each applicant with post-effort rank  $\theta_{\text{post}} \geq c$  is admitted with probability  $\ell_1 = \frac{\rho}{1-c} \in (0, 1]$ . Others are rejected,  $\ell_0 = 0$ .*

Note that standard non-randomized admissions policies are equivalent to case where  $c = 1 - \rho$ : the highest scoring applicants up to the capacity constraint are accepted with probability 1 and all others are rejected. We call this case **non-randomized admissions**. The other extreme is the one-level policy, **pure randomization**, where each applicant is admitted with probability  $\rho$ , i.e.,  $\ell_0 = \rho$ . Decreasing the cut-off  $c$  can be viewed as increasing the level of randomization in the admissions policy.

We now reason about how various welfare and fairness metrics vary with  $\lambda$  (in the two-level policy class, just  $c$ ). To simplify the presentation, we assume that  $f, g$  and  $p$  are differentiable and that baseline effort  $e_0$  is 0.

**Applicant welfare** The following result shows how overall welfare is maximized.

<sup>6</sup> These were:  $f(x) = 2x$  (corresponding to a uniform distribution of skill levels),  $g(x) = \sqrt{x}$  (decreasing marginal returns to effort) and  $p = x^2$  (increasing marginal cost of effort).

<sup>7</sup>See Corollary A.1 for an example of how equilibria change with the function  $\lambda$

**Proposition 3.1** (Applicant welfare). *Among all  $\lambda$  with  $K$  levels, for  $K \geq 1$ , applicant welfare  $\mathcal{W}$  is maximized by the one-level policy with pure randomization. Further, in the class of two-level policies,  $\mathcal{W}$  is monotonically non-increasing in  $c$ .*

This result is perhaps unsurprising given our formulation of applicant welfare – in which effort is costly but applicants do not directly benefit from their score  $v$ , only their resulting admissions probabilities  $\lambda(\theta_{\text{post}})$ . Thus,  $\mathcal{W}$  is maximized by a completely random admissions policy where the cost of exerting effort is  $p(e_0) = 0$  for every applicant.

**School’s private utility** On the other extreme is the school’s private utility  $\mathcal{U}^{\text{pri}}$ : in which only the expected scores of admitted applicants matters. For  $K = 2$ , it is maximized by a deterministic decision policy.

**Proposition 3.2** (School’s private utility for two-level policies). *In a two-level policy, the school’s private utility  $\mathcal{U}^{\text{pri}}$  is monotonically non-decreasing in  $c$ , and consequently is maximized by non-randomized admissions.*

The result follows because both aspects of the school’s private utility increases with the cutoff  $c$  in a two-level policy: first, as  $c$  increases, the school admits higher ranked applicants and the scores of the admitted applicants increase with the rank of the admitted applicants; second, increasing  $c$  increases  $\ell_1$ , which, as per Corollary A.1, further increases the equilibrium effort levels (and thus the post-effort scores) of the highest ranked applicants. Note that this effect occurs even though rankings of applicants are identical under any  $\lambda$  (Proposition 2.1).

A natural question is whether the deterministic decision policy also maximizes  $\mathcal{U}^{\text{pri}}$  among all  $K$ -level policies for  $K > 2$ . Surprisingly, the answer in general is negative.

**Proposition 3.3.** *A 3-level policy may achieve strictly higher  $\mathcal{U}^{\text{pri}}$  than non-randomized admissions.*

This result provides a perhaps counter-intuitive insight for practice: *even for a school maximizing its own utility*, deterministically accepting the top students is not generally optimal – schools could improve on both student welfare and its own utility by randomizing. The counter-example used to prove the above proposition involves picking a skill distribution (the distribution of  $f(\theta_{\text{true}})$ ) with a long tail. That is, if the skill level  $f(\theta_{\text{pre}})$  of applicants with rankings above the optimal deterministic cutoff  $1 - \rho$  is relatively high, then the school can improve  $\mathcal{U}^{\text{pri}}$  by using a three-level policy that increases the competition for high admission probability and incentivizes higher scores among the top ranked applicants.

**Societal utility** The previous results show that for both applicant welfare and school’s private utility, an extreme two-level function is optimal among all two-level admissions functions  $\lambda$ : pure randomization for applicant welfare, and non-randomization for school’s private utility. Next, we show that a similar result does not hold for societal utility,  $U^{\text{soc}} = \mathbb{E}[v]$ : if the goal is to maximize the score attained by the entire applicant population and not just the admitted applicants, the admissions function  $\lambda$  should be randomized (but not purely randomized), even among two-level policies.<sup>8</sup>

**Proposition 3.4** (Societal utility maximization for two-level policies). *Among two-level functions, there exists a setting in which societal utility  $U^{\text{soc}}$  is maximized by choosing  $c \in (0, 1 - \rho)$ ,  $\ell_1 = \frac{\rho}{1-c} \in (\rho, 1)$ .*

Here, increasing  $c$  and  $\ell_1$  has two competing effects: (1) it increases the effort level of the highest ranked applicants in  $\psi_1$ ; (2) it simultaneously increases the fraction of applicants who exert minimum effort, as they go from being admitted with probability  $\ell_1$  (and thus needing to exert effort to preserve their position) to being deterministically rejected. Maximizing overall score across the population thus requires an intermediate level of randomization. More generally, a  $K$ -level reward function for  $K > 2$  may improve societal utility over a two-level policy.

The choice of  $\lambda$  not only differentially affects applicants’ admissions probabilities; it also differentially affects applicant *efforts* and *scores* in equilibrium. Choosing a reward function  $\lambda$  to increase the equilibrium scores of some applicants comes at a cost of the scores of other applicants. Thus, as the contrast between  $U^{\text{soc}}$  and  $U^{\text{pri}}$  establishes, the optimal design depends on which applicants one considers. When a school has the power to choose design  $\lambda$  to maximize its own private interests, doing so comes at a cost of scores of other applicants. As shown in Figure 2, the exact tradeoffs between  $U^{\text{soc}}$ ,  $U^{\text{pri}}$ , and applicant welfare  $\mathcal{W}$  depend on various model parameters, such as  $f$ ,  $g$  and  $p$ , as well as the capacity  $\rho$ .

## 4 ENVIRONMENT DIFFERENCES

We have thus far assumed a basic model where applicants differ only in their latent skill levels  $f(\theta_{\text{pre}})$ , and their observed score depends only on their latent skill level and chosen effort. However, in societies

<sup>8</sup>Our notion of societal utility differs from the traditional notion of *social welfare*, which would also include in its formulation the effort costs of applicants. Proposition 3.1 and Proposition 3.4 would imply that social welfare is also maximized by an intermediate level of randomization.




Figure 2:  $U^{\text{soc}}$ ,  $U^{\text{pri}}$ , and  $\mathcal{W}$  for two-level policies parametrized by  $c \in (0, 1 - \rho)$ . Model parameters are as in Figure 1. The school’s capacity is  $\rho = 0.2$ .

with structural inequalities, an individual’s measured success also depends on various environmental factors beyond one’s control, such as family income and the availability of resources in one’s community. Roemer (1998)’s work on the equality of opportunity argues that environmental factors (called “circumstances”) are distinguished from effort. For example, for the same amount of time spent studying, an applicant from a well-resourced school may achieve higher test scores than one from an under-resourced school. In this section, we extend the model introduced in Section 2 to study the disparate impact of admission policies in the presence of structural inequalities, specifically differences in the applicant’s previous education environment.

**Model and equilibria characterization** We now denote each applicant’s latent skill rank as  $\theta_{\text{true}} \in [0, 1]$ . In addition to the latent skill, each applicant has an (unobserved) environmental factor  $\psi \in \Psi$  that represents how favorable their environment is for attaining a higher score. Because a favorable environment results in a higher rate of return for effort, we model the environment as a multiplicative factor in the score (see e.g., Calsamiglia (2009)). Formally, the post-effort score is a function of the latent skill, the environmental factor, and the effort level:

$$v = \psi \cdot g(e) \cdot f(\theta_{\text{true}}),$$

where  $g, f$  are as defined in Section 2.1.

We assume there are two groups of applicants, A and B, and the distribution of skill is the same in both groups. Group A has a more favorable environment factor, that is,  $\Psi = \{\psi_A, \psi_B\}$  and  $\psi_A > \psi_B$ . Thus

we will also refer to  $\mathbf{B}$  as the “disadvantaged group”. To simplify our presentation, we assume each group is half of the total applicant population, though the results in this section generalize. We defer all proofs in this section to Appendix F.

We begin by characterizing the equilibrium ranking under the designer’s policy  $\lambda$ . The equilibrium effort levels and post-effort ranks are as defined in Definition 1, except they are now group-dependent, that is, we have  $e(\theta_{\text{true}}, \psi)$  and  $\theta_{\text{post}}(\theta_{\text{true}}, \psi)$ . Because of the differences in  $\psi$ , the post-effort ranking  $\theta_{\text{post}}$  is now group-dependent, and in general is not equal to  $\theta_{\text{true}}$ . To apply Proposition 2.1 as before, we construct an “environment-scaled pre-effort rank” (denoted  $\theta_{\text{pre}}$ ).

**Proposition 4.1** (Equilibrium under group differences). *Define  $\theta_{\text{pre}}$  as:  $\theta_{\text{pre}} := f_{\text{mix}}^{-1}(f(\theta_{\text{true}}) \cdot \psi)$ , where  $f_{\text{mix}}^{-1}$  is the CDF for the environment-scaled skill,  $f(\theta_{\text{true}}) \cdot \psi$ :*

$$f_{\text{mix}}^{-1}(x) := \frac{1}{2}f^{-1}(x/\psi_{\mathbf{A}}) + \frac{1}{2}f^{-1}(x/\psi_{\mathbf{B}}).$$

*Then, in every equilibrium,  $\lambda(\theta_{\text{post}}(\theta_{\text{true}}, \psi)) = \lambda(\theta_{\text{pre}})$ .*

The environment factor depresses the  $\theta_{\text{pre}}$ , and therefore  $\theta_{\text{post}}$ , of the disadvantaged group. Rank preservation (Proposition 2.1) again simplifies equilibrium analysis.

**Welfare gap and access differences** Using the above characterization, we now study the disparate impact of various rank reward functions  $\lambda$ , through two key facets: (1) the *welfare gap* and (2) *access*. The former measures welfare disparity between applicants from different groups *with the same latent skill rank*.

**Definition 3** (Welfare gap). *Let  $\mathcal{W}^{\mathbf{G}}(\theta_{\text{true}})$  denote post-effort welfare of an applicant with latent skill ranking  $\theta_{\text{true}}$  from group  $\mathbf{G} \in \{\mathbf{A}, \mathbf{B}\}$ , i.e.,*

$$\mathcal{W}^{\mathbf{G}}(\theta_{\text{true}}) := \lambda(\theta_{\text{post}}(\theta_{\text{true}}, \psi_{\mathbf{G}})) - p(e(\theta_{\text{true}}, \psi_{\mathbf{G}})). \quad (1)$$

*We define the welfare gap as*

$$\mathcal{G}(\theta_{\text{true}}) := \mathcal{W}^{\mathbf{A}}(\theta_{\text{true}}) - \mathcal{W}^{\mathbf{B}}(\theta_{\text{true}}).$$

The welfare gap captures differences in admission probabilities *and* in the effort needed to achieve such probabilities. Our next notion, *access*, captures whether a decision policy *includes* the disadvantaged group in the admitted class, regardless of effort.

**Definition 4** (Access). *Access is the overall probability of admission of the disadvantaged group.*

$$\mathcal{A} := \mathbb{E}_{\theta_{\text{true}}}[\lambda(\theta_{\text{post}}(\theta_{\text{true}}, \psi_{\mathbf{B}}))].$$

**Proposition 4.2** (Admission and pointwise welfare gap for two-level policies). *Denote the group-specific rank threshold for group  $\mathbf{G}$  as*

$$\text{thres}_{\mathbf{G}}(c) := f^{-1}\left(\frac{f_{\text{mix}}(c)}{\psi_{\mathbf{G}}}\right).$$

*A two-level policy with  $c \in (0, 1 - \rho]$  admits a group  $\mathbf{A}$  applicant with  $\theta_{\text{true}} \geq \text{thres}_{\mathbf{A}}(c)$  with probability  $\frac{\rho}{1-c}$ , a group  $\mathbf{B}$  applicant with  $\theta_{\text{true}} \geq \text{thres}_{\mathbf{B}}(c)$  with probability  $\frac{\rho}{1-c}$ , and all other applicants with probability 0. The welfare gap  $\mathcal{G}(\theta_{\text{true}})$  is non-negative for every  $\theta_{\text{true}}$ , and strictly positive for  $\theta_{\text{true}} \geq \text{thres}_{\mathbf{B}}(c)$ . In contrast, the one-level pure randomization policy has  $\mathcal{G}(\theta_{\text{true}}) \equiv 0$ .*

In the “High” region of  $\theta_{\text{true}}$  (where  $\theta_{\text{true}} \geq \text{thres}_{\mathbf{B}}(c)$ ), both the group  $\mathbf{A}$  and group  $\mathbf{B}$  applicant have the same probability of admission. However, group  $\mathbf{B}$  applicants must expend more effort, resulting in a strictly positive welfare gap. In the “Middle” region (where  $\theta_{\text{true}} \in [\text{thres}_{\mathbf{A}}(c), \text{thres}_{\mathbf{B}}(c))$ ), group  $\mathbf{A}$  applicants are admitted with positive probability while group  $\mathbf{B}$  applicants are deterministically rejected, again leading to a positive welfare gap. In the “Low” region (where  $\theta_{\text{true}} < \text{thres}_{\mathbf{A}}(c)$ ), applicants from both groups are deterministically rejected, leading to no gap. Further comparison of the equilibrium welfare, admission probability and effort for any given two-level policy can be found in Table 1 in the appendix.

The previous proposition highlights the joint role of admissions probability and effort in determining the welfare gap. Our next result focuses on the welfare gap in the “High” region, where applicants of both groups are admitted with the same positive probability, showing that the welfare gap decreases as a two-level admission policy becomes more randomized.

**Proposition 4.3** (Welfare gap increases with  $c$ ). *Consider the setting in Definition 2, with the school’s chosen admissions policy  $c = \bar{c}$ , where  $\bar{c} \leq (f_{\text{mix}})^{-1}(f(1) \cdot \psi_{\mathbf{B}})$ . Then, for any  $\theta_{\text{true}} \geq \text{thres}_{\mathbf{B}}(\bar{c})$ , that is,  $\theta_{\text{true}}$  is in the “High” region, we have  $\left.\frac{\partial \mathcal{G}(\theta_{\text{true}})}{\partial c}\right|_{c=\bar{c}} > 0$ .*

Decreasing  $c$  (increasing randomization) reduces the welfare gap in the “High” region of  $\theta_{\text{true}}$ . As the school’s private utility is actually increasing in  $c$  (Proposition 3.2), there is also a tradeoff between the school’s private utility and the welfare gap.

Increasing randomization also increases access: it is maximized by pure randomization and there exists a large class of models where access is always improved by more randomization in a two-level policy.

**Proposition 4.4** (Access decreases with  $c$ ). *Pure randomization has higher  $\mathcal{A}$  than any two-level policy.*



Moreover, if  $f^{-1}$  is convex,  $\mathcal{A}$  for two-level policies is non-increasing in  $c$ .

By reducing the level of competition among applicants, the more randomized admission policy increases access, becoming more inclusive to applicants who are disadvantaged by their environment. A ranking reward function that increases the incentive for applicants to compete tends to increase the welfare gap and reduce access. Without knowledge of the applicant’s environment factor, randomizing the admissions policy thus may be key to reducing disparate outcomes.

## 5 DISCUSSION

We now draw some comparisons to welfare results known in strategic classification. Taken together, Proposition 3.1 and 3.2 imply, among two-level policies, that there is a direct tradeoff between the applicant welfare and the school’s private utility: by reducing the degree of randomization in the admission policy, we increase private utility at the expense of applicant welfare. In the strategic classification setting, Milli et al. (2019) observed a similar tradeoff between the “institution utility” (the classification loss under gaming) and the “social burden” (cost incurred by individuals for changing their features). Both sets of results suggest that strategic behavior complicates the choice of an optimal decision policy and necessitates careful adjudication among stakeholders.

Our key takeaway that randomization in the decision policy  $\lambda$  can improve applicant welfare is also related to Braverman and Garg (2020). They studied the welfare benefits of randomization—in the form of probabilistic classifiers and noisy features—and observed that the designer has no incentive to use a more randomized classifier. This finding is true in our setting, only for the class of two-level policies – randomization with more levels may improve both welfare and private school utility.

The strategic classification literature (Braverman and Garg, 2020; Hu et al., 2019; Milli et al., 2019) also considers the disparate costs of strategic behavior. These works study two groups with different costs of gaming, and Milli et al. (2019) introduces the concept of “social gap”: the difference between the costs of successful gaming incurred. Our *welfare gap* can be seen as a measure of the difference in the costs of effort in the “High” region of latent skill rank. We showed a tradeoff between the school’s private utility and the welfare gap (Proposition 3.2 and 4.3), while Milli et al. (2019) (Theorem 3.1) showed a tradeoff between the institutional utility and the social gap. On the other hand, the concept of access—the proportion of the dis-

advantaged group admitted—as an indicator of disparate impact has not received much attention in the strategic classification literature, as one cannot compare such access across mechanisms that accept different numbers of people overall.

Overall, we believe that our strategic ranking model is a natural one through which to study constrained allocation settings for strategic decision making, such as admissions and hiring. While our base model is general, equity notions in other applications may differ from those considered in Section 4. From a technical perspective, there remain open computational questions for finding optimal  $K$  level reward functions. More generally, our work supports the consideration of the ranking and constrained allocation model over unconstrained classification in the FATE ML community, centering the study of welfare notions of fairness over that of statistical parity.

## Acknowledgements

The authors thank Frances Ding and John Miller for helpful comments on a draft. They also thank the anonymous reviewers for their invaluable feedback.

## References

- R. Abebe and K. Goldner. Mechanism design for social good. *AI Matters*, 4(3):27–34, 2018. 3
- T. Alon, M. Dobson, A. Procaccia, I. Talgam-Cohen, and J. Tucker-Foltz. Multiagent evaluation mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1774–1781, 2020. 2
- H. Aziz. Developments in multi-agent fair allocation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13563–13568, 2020. 3
- Y. Barut and D. Kovenock. The symmetric multiple prize all-pay auction with complete information. *European Journal of Political Economy*, 14(4):627–644, 1998. 2
- Y. Bechavod, K. Ligett, S. Wu, and J. Ziani. Gaming helps! learning from strategic interactions in natural dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 1234–1242. PMLR, 2021. 2
- G. S. Becker. A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846, 1973. 3
- A. L. Bodoh-Creed and B. R. Hickman. College assignment as a large contest. *Journal of Economic Theory*, 175:88–126, 2018. 1, 2, 6, 13
- M. Braverman and S. Garg. The role of randomness and noise in strategic classification. *arXiv preprint arXiv:2005.08377*, 2020. 2, 9, 13

- M. Brückner, C. Kanzow, and T. Scheffer. Static prediction games for adversarial learning problems. *The Journal of Machine Learning Research*, 13(1):2617–2654, 2012. 1, 2
- W. Cai, J. Gaebler, N. Garg, and S. Goel. Fair allocation through selective information acquisition. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, page 2228, New York, NY, USA, 2020. Association for Computing Machinery. 3
- C. Calsamiglia. Decentralizing equality of opportunity\*. *International Economic Review*, 50(1):273–290, 2009. 3, 7
- A. Chouldechova and A. Roth. The Frontiers of Fairness in Machine Learning. *CoRR*, abs/1810.08810, 2018. 3
- B. L. Connelly, L. Tihanyi, T. R. Crook, and K. A. Gangloff. Tournament theory: Thirty years of contests and competitions. *Journal of Management*, 40(1):16–47, 2014. 2
- L. C. Corchón. The theory of contests: a survey. *Review of economic design*, 11(2):69–100, 2007. 2
- J. Dong, A. Roth, Z. Schutzman, B. Waggoner, and Z. S. Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018. 2
- E. Elkind, A. Ghosh, and P. Goldberg. Contest design with threshold objectives, 2021. 3
- Y. Faenza, S. Gupta, and X. Zhang. Impact of bias on school admissions and targeted interventions. *arXiv preprint arXiv:2004.10846*, 2020. 3
- D. Fang, T. Noe, and P. Strack. Turning up the heat: The discouraging effect of competition in contests. *Journal of Political Economy*, 128(5):1940–1975, 2020. 2
- Q. Fu and Z. Wu. Contests: Theory and topics. In *Oxford Research Encyclopedia of Economics and Finance*. 2019. 2
- N. Garg, H. Li, and F. Monachou. Dropping standardized testing for admissions: Differential variance and access. *arXiv preprint arXiv:2010.04396*, 2020. 3
- P. Gözl, A. Kahng, and A. D. Procaccia. Paradoxes in fair machine learning. *NeurIPS'19*, 2019. 3
- N. Haghtalab, N. Immorlica, B. Lucier, and J. Z. Wang. Maximizing welfare with incentive-aware evaluation mechanisms. *arXiv preprint arXiv:2011.01956*, 2020. 2
- E. A. Hanushek and L. Woessmann. *The high cost of low educational performance: The long-run economic impact of improving PISA outcomes*. ERIC, 2010. 4
- M. Hardt, N. Megiddo, C. Papadimitriou, and M. Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, pages 111–122, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4057-1. 1, 2
- B. Holmstrom and P. Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.*, 7:24, 1991. 16, 17
- L. Hu and Y. Chen. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 1389–1398, Republic and Canton of Geneva, Switzerland, 2018. International World Wide Web Conferences Steering Committee. ISBN 978-1-4503-5639-8. 3
- L. Hu, N. Immorlica, and J. W. Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT\* '19, pages 259–268, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6125-5. doi: 10.1145/3287560.3287597. 2, 9
- N. Immorlica, K. Ligett, and J. Ziani. Access to population-level signaling as a source of inequality. In *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*, pages 249–258, 2019. 3
- S. Kannan, A. Roth, and J. Ziani. Downstream effects of affirmative action. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 240–248, 2019. 3
- S. Kannan, M. Niu, A. Roth, and R. Vohra. Best vs. all: Equity and accuracy of standardized test score reporting. *arXiv preprint arXiv:2102.07809*, 2021. 3
- J. Kleinberg and M. Raghavan. How Do Classifiers Induce Agents to Invest Effort Strategically? In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, pages 825–844, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6792-9. 2, 15
- K. A. Konrad. Strategy in contests—an introduction. *WZB-Markets and Politics Working Paper No. SP II*, 1, 2007. 2
- I. Lindenlaub and F. Postel-Vinay. Multidimensional sorting under random search. *Manuscript, University College London*, 2016. 15
- L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pages 3150–3158. PMLR, 2018. 3

- L. T. Liu, A. Wilson, N. Haghtalab, A. T. Kalai, C. Borgs, and J. Chayes. The disparate equilibria of algorithmic decision making when individuals invest rationally. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 381–391, 2020. 3
- Z. Liu and N. Garg. Test-optional policies: Overcoming strategic behavior and informational gaps. *arXiv preprint arXiv:2107.08922*, 2021. 3
- M. Mathioudakis, C. Castillo, G. Barnabo, and S. Celis. Affirmative action policies for top-k candidates selection: with an application to the design of policies for university admissions. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, pages 440–449, 2020. 3
- J. Miller, S. Milli, and M. Hardt. Strategic classification is causal modeling in disguise. In *International Conference on Machine Learning*, pages 6917–6926. PMLR, 2020. 2
- S. Milli, J. Miller, A. D. Dragan, and M. Hardt. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, pages 230–239, New York, NY, USA, 2019. ACM. 2, 9
- H. Mouzannar, M. I. Ohanessian, and N. Srebro. From fair decision making to social equality. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, pages 359–368, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6125-5. 3
- R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. 5, 13
- A. Noriega-Campero, M. A. Bakker, B. Garcia-Bulle, and A. Pentland. Active fairness in algorithmic decision making. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 77–83, 2019. 3
- W. Olszewski and R. Siegel. Large contests. *Econometrica*, 84(2):835–854, 2016. 1, 2, 3, 6, 13
- W. Olszewski and R. Siegel. Bid caps in large contests. *Games and Economic Behavior*, 115:101–112, 2019a. 1, 2, 3
- W. Olszewski and R. Siegel. Pareto improvements in the contest for college admissions. *Unpublished paper, Department of Economics, Northwestern University.[1607]*, 2019b. 1, 2, 3, 13
- J. E. Roemer. *Equality of opportunity*. Harvard University Press, 1998. 3, 7
- E. Rolf, M. Simchowitz, S. Dean, L. T. Liu, D. Bjorkegren, M. Hardt, and J. Blumensstock. Balancing competing objectives with noisy data: Score-based classifiers for welfare-aware machine learning. In *International Conference on Machine Learning*, pages 8158–8168. PMLR, 2020. 3
- D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, Apr 1973. 3
- Y. Shavit, B. Edelman, and B. Axelrod. Causal strategic linear regression. In *International Conference on Machine Learning*, pages 8676–8686. PMLR, 2020. 2
- M. Spence. Job market signaling. In *Uncertainty in economics*, pages 281–306. Elsevier, 1978. 2, 3
- B. Tabibian, V. Gomez, A. De, B. Schölkopf, and M. G. Rodriguez. On the design of consequential ranking algorithms. In *Conference on Uncertainty in Artificial Intelligence*, pages 171–180. PMLR, 2020. 3
- M. Zehlike and C. Castillo. Reducing disparate exposure in ranking: A learning to rank approach. In *Proceedings of The Web Conference 2020*, pages 2849–2855, 2020. 3
- M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. Baeza-Yates. Fa\* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1569–1578, 2017. 3

---

**Supplementary Material:**  
**Strategic ranking**

---

## A Supplemental model discussion

Here we include supplemental technical remarks on the model.

**Remark A.1** (Construction of the  $\gamma$  map). *Formally, if  $v_0$  is an atom of the score distribution and  $\Omega_{v_0}$  is the set of tied applicants with score  $v_0$ , then the discontinuity of  $F$  at  $v_0$  has height equal to the measure of  $\Omega_{v_0}$ , and so  $\gamma$  for those applicants can be filled by using the CDF of  $\Gamma(\omega)$  restricted to  $\Omega_{v_0}$ ; this gives a distribution for  $\gamma(\omega, v_0)$  that is uniform over the gap interval  $[F_-(v_0), F(v_0)]$  when restricted to  $\Omega_{v_0}$ , and hence leads to the claimed uniform distribution of  $\theta_{\text{post}}(\omega)$ . Finally, we define  $\gamma(\omega, v)$  for applicants  $\omega \notin \Omega_{v_0}$  by “slotting them in” in such a way that for a pair  $\omega, \omega'$  with exactly one member in  $\Omega_{v_0}$  and  $\Gamma(\omega) < \Gamma(\omega')$  we have that  $\gamma(v_0, \omega) \leq \gamma(v_0, \omega')$ . Since the image of  $\Omega_{v_0}$  is by construction dense in the gap interval, this uniquely determines  $\gamma(v_0, \omega)$  for all applicants  $\omega$ .*

*Note that for a given applicant  $\omega$ , the map  $v \mapsto \gamma(\omega, v)$  is not necessarily 1-1; indeed, if  $F$  is constant on an interval  $I$ , all  $v \in I$  lead to the same rank. But this only effects regions where the distribution of  $v$  has no mass, and thus will not cause any issues.*

**Remark A.2** (Ties in pre-effort skill). *With probability one, there are no two applicants with the same skill level, and the support of the skill distribution has no gap: given two applicants with different skills, the probability of finding an applicant with skill in between these two is always non-zero.*

**Remark A.3** ( $\gamma$ 's dependency on the set of efforts). *Formally,  $\gamma$  depends on the set of efforts. Given a fixed set and  $\gamma$ , for each applicant  $\omega$  the first condition considers the counter-factual ranking of  $\omega$  with different post-effort values but using the same ranking function. As defined,  $\gamma$  yields a uniform distribution of ranks with such measure 0 changes. In Appendix Lemma C.2, we further prove that two effort sets equal up to sets of measure 0 induce the same ranking function  $\gamma$ , and so the condition is consistent.*

**Remark A.4** (Credible commitment on the part of the school). *Note that we do not require that, given applicant effort levels, the design  $\lambda$  is optimal for the school's utility. As in e.g. Braverman and Garg (2020) for classification, the randomized ranking reward (characterized by probability  $\lambda$ ) is not optimal for the school after applicants have chosen their effort levels in response to the classifier. Such a characterization of equilibria thus requires credible commitment on the part of the school.*

**Remark A.5** (Interpreting  $\mathcal{U}^{\text{pri}}$  as a conditional expectation.). *In our primary admissions interpretation,  $\lambda(\theta_{\text{post}})$  is a probability. (All our results also hold when  $\lambda(\theta_{\text{post}})$  represents a deterministic reward.) In a setting where there are a finite number of applicants (as opposed to our continuum model) and the admission outcome of each applicant with rank  $\theta_{\text{pre}}$  is  $Z \sim \text{Bernoulli}(\lambda(\theta_{\text{post}}))$ ,  $\mathcal{U}^{\text{pri}}$  can also be interpreted as a conditional expectation (average score of admitted applicants), that is  $\mathbb{E}[\mathbb{E}[v \mid Z = 1]]$ .*

**Remark A.6** (Comparison to related work in contests literature). *Proposition 2.1 has analogues in the contests literature, in particular Bodoh-Creed and Hickman (2018) and Olszewski and Siegel (2016). Here we remark on the technical differences with these results. Bodoh-Creed and Hickman (2018) assumes that the density of prizes has full support—this assumption, translated to our setting, requires the ranking reward function  $\lambda$  to be continuous. In contrast, Proposition 2.1 is proven for  $\lambda$  that is a discontinuous step-function. Theorem 2(a) of Olszewski and Siegel (2016), while not requiring  $\lambda$  to be continuous, holds for all but a small fraction of the applicants. In other words, Proposition 2.1, though recalls the assortative allocation principle known to the contests literature, is neither a restatement of previous results nor a generalization. It is worthwhile to note that we prove Proposition 2.1 using elementary arguments that may be of independent interest.*

*Olszewski and Siegel (2016, 2019b) also remark that the unique mechanism that implements assortative allocation is given by Myerson (1981), assuming quasi-linear utility. However, due to the non-quasi-linearity of our utility function  $U$  in the score  $v(e, \theta_{\text{pre}})$ , the equilibrium strategy in our setting (Theorem 2.2) turns out to be a variant of Myerson's payment rule, and is derived from an independent analysis.*

**Corollary A.1** (Effort comparative statics). *Assume that  $g$  and  $p$  are differentiable. If  $\ell_k$  increases and  $\ell_j$  decreases for some  $j > k$  (fixing all other parameters), then  $e_i(\theta)$  for all  $i < k$  are unaffected,  $e_k$  is weakly increasing, and efforts  $e_{k+1}, \dots, e_j$  are weakly decreasing.*

Perhaps surprisingly, increasing a reward  $\ell_k$  does not affect the equilibrium effort of the applicants in the band  $k-1$  immediately below: in equilibrium, they do not receive the higher reward since those in band  $k$  correspondingly increase their effort. The proof of the corollary actually implies that  $e_k, e_{k+1}$  and  $e_j$  are strictly monotonic in the part of  $\psi_k, \psi_{k+1}$  and  $\psi_j$  where the efforts are strictly above  $e_0$  – which depending on the parameters of

the model can just be part of these intervals, or all of these intervals, as observed in the paragraph following Theorem 2.2.

## B Multi-dimensional skill

Up to now, we have studied the setting in which there is a single-dimensional measurable score on which applicants are ranked and can exert effort. In this section we consider an extension of our model to  $m$  skills, and use the extension to explore two questions of interest: (1) How does the school’s reward design influence the applicant’s decision to allocate effort across different skills in the competitive setting? (2) When one of the skills is valued but *not* measurable by the school, how does competition in the measurable skill trade-off affect the school’s utility? We first introduce the model extension, and then present results in Sections B.1 and B.2 addressing questions (1) and (2) respectively. We note that the results in this section are preliminary; we end the section with a discussion on directions for future inquiry. All proofs in this section can be found in Appendix G.

**Model.** The model is similar to our base model Section 2.1; each applicant now has  $m$  latent skill levels with respective ranks  $\theta_{\text{pre}}^i \in [0, 1]$  for  $i \in \{1, \dots, m\}$ . Each rank is drawn independently from the uniform distribution over  $[0, 1]$ , and the skill  $i$  of applicant with rank  $\theta_{\text{pre}}^i$  is  $f_i(\theta_{\text{pre}}^i)$ . Each applicant now chooses effort levels  $\{e_i\}$ , at cost  $p^m(e_1, \dots, e_m)$ , resulting in post-effort scores  $\{v_i\}$ ,

$$v_i = g(e_i) \cdot f_i(\theta_{\text{pre}}^i),$$

where  $g$  is concave, increasing, as before, and  $f_i$  is a continuous, strictly increasing quantile function ( $f_i^{-1}$  is the CDF function of the scores on dimension  $i$ ). As before, the school observes the post-effort scores for each applicant, now for each dimension  $i$ , and designs a non-decreasing function  $\lambda : [0, 1] \rightarrow [0, 1]$  denoting the admissions probability  $\lambda(\theta_{\text{post}})$  for applicant with post-effort rank  $\theta_{\text{post}}$ .

How does the school construct post-effort rank  $\theta_{\text{post}}$ ? In general, each applicant’s *combined* post-effort score may be any function of the scores on each dimension,  $\{v_i\}$ . Here, we assume the following linear score function. The school announces weights  $\alpha = (\alpha_1, \dots, \alpha_m) \in \Delta_{m-1}$  (denoting the  $m$ -dimensional simplex); the weights  $\alpha$  represent the relative emphasis placed on each skill for the admissions decision. Then, each applicant is ranked according to their combined score,  $v^\alpha := \sum_{i=1}^m \alpha_i v_i$ , resulting in their combined post-effort rank  $\theta_{\text{post}}^\alpha$ , and is admitted with probability  $\lambda(\theta_{\text{post}}^\alpha)$ .<sup>9</sup>

Putting things together, the applicant’s individual welfare is:

$$W(\{e_i\}_{i=1}^m, \lambda(\theta_{\text{post}}^\alpha)) = \lambda(\theta_{\text{post}}^\alpha) - p^m(e_1, \dots, e_m)$$

### B.1 Equilibrium under multi-dimensional competition

In this section, we apply the model for multi-dimensional skills described above to study equilibrium effort allocations and rank preservation. In the single skill case, we found that the post-effort ranks equaled the pre-effort ranks (the allocation preserves rank) (Proposition 2.1). As we’ll see, post-effort rank preservation is more complex with a multi-dimensional score. In fact, to our knowledge, the economics contests literature has not considered multi-dimensional scores, and defining assortative allocations (the analogue of rank preservation) in other domains (such as participant search in a matching market setting) has proven tricky (Lindenlaub and Postel-Vinay, 2016). Our first result for the multi-dimensional case is that rank preservation no longer holds if pre-effort ranks are defined as each applicant’s pre-effort skill combination,  $\alpha \cdot f(\theta_{\text{pre}})$ . Rather, the post-effort ranks may depend on the weights  $\alpha$  and the distribution of the skill levels,  $f_i$ ; however, it may be possible to define an alternative pre-effort rank function under which rankings are preserved.

We show this result by characterizing the equilibrium of the following simplified setting, with further assumptions on the effort cost function  $p$  and the effort transfer function  $g$ ; the cost function  $p$  is assumed to be a function of the sum of the efforts exerted:

$$p(e^1, \dots, e^m) \triangleq p\left(\sum_{i=1}^m e_i\right),$$

where  $p$  is convex and increasing. This assumption says that the effort exerted for any skill is entirely exchangeable, for example, two hours spent on studying math is as costly as two hours spent on studying chemistry. Effort transfer function  $g$  is assumed to be linear, that is, there are constant returns to effort.

<sup>9</sup>The linear combination of skill is similar to the “linear mechanism” in Kleinberg and Raghavan (2019), which studied how reward design incentivizes strategic agents to exert effort on different skill dimensions. Compared to their work, where efforts are connected to skills via an effort graph, we consider a simplified setting where each effort maps to one skill, and study how reward design affects equilibrium rankings, in presence of competition.

Under these assumption, applicants are incentivized to put effort into a single skill; rankings are preserved not on the pre-effort skill combination but rather just their most important skill dimension.

**Proposition B.1** (Multi-dimensional rank preservation for linear  $g$ ). *Suppose  $g$  is a linear function such that  $g(e) = hx$ ,  $h > 0$ . Suppose the school picks some  $\alpha$  and  $\lambda$ . Define the combined pre-effort index as:*

$$v_{pre}^\alpha := \max_i \alpha_i f_i(\theta_{pre}^i).$$

*Then in every equilibrium, for any two applicants with combined pre-effort indices  $v_{pre}^\alpha, \overline{v_{pre}^\alpha}$  and combined post-effort ranks  $\theta_{post}^\alpha, \overline{\theta_{post}^\alpha}$  we have*

$$\lambda(\theta_{post}^\alpha) > \lambda(\overline{\theta_{post}^\alpha}) \iff v_{pre}^\alpha > \overline{v_{pre}^\alpha}.$$

Each applicant's decision to exert effort at equilibrium now also depends on the skill-specific quantile  $f_i$ , whereas in the single skill setting, only the pre-effort rank was relevant. Applicants who have a skill that they have a large advantage in *relative* to the rest of the applicants (high  $f_i(\theta_{pre}^i)$ ) and that is valued highly by the school (high  $\alpha_i$ ) are advantaged in terms of post-effort rankings.

For non-linear  $g$ , there may not exist a simple characterization of the post-effort ranks. For example, when  $g$  is strictly concave, that is, there are decreasing marginal returns to effort for each skill, it is no longer optimal for a applicant to only put all their effort in one skill, and the relative allocation of effort will depend on the specific functional form of  $g$ .

## B.2 School's private utility with unmeasurable skill

We now analyze a multi-dimensional setting where one skill dimension is *unobservable* though still valuable to the school. This setting is directly motivated by the classic contract design work of [Holmstrom and Milgrom \(1991\)](#), who show that when some work tasks are less measurable than others, it may be optimal to lessen incentives on the measurable tasks (e.g., by adopting a fixed wage) so as to not crowd out effort in the less measurable ones.

In our simplified setting, there are two skills  $M$  and  $U$ :  $M$  has a measurable score  $v^M$  and  $U$  has an unmeasurable score  $v^U$ . For example,  $M$  could be scholastic achievement as measured by SAT scores, and  $U$  could be "creativity", a personal quality that is valued by the school but is not directly measurable. Since the school cannot observe  $v^U$ , its admission policy is based on  $\theta_{post}^M$  only, that is,  $\alpha_U = 0$  and  $\alpha_M = 1$ .

We further assume that each applicant has a fixed effort budget of  $B > 0$ , and is intrinsically motivated to exert effort in the unmeasurable skill  $U$ ; in fact, they will always exert effort  $e^U = B - e^M$ . Formally this corresponds to the effort cost function

$$p^m(e_M, e_U) = p(e^M) - (\max(0, B - (e^M + e^U)))^2$$

where  $p$  is convex and increasing. We can write the applicant's individual welfare as:

$$W(e^M, e^U, \lambda(\theta_{post}^M)) = \lambda(\theta_{post}^M) - p^m(e_M, e_U).$$

The school's private utility is now weighted by  $\beta \in (0, 1)$ , which quantifies the relative value the school places on the measurable skill over the unmeasurable skill.

$$\mathcal{U}_\beta^{\text{pri}} = \mathbb{E}[\beta \cdot v^M + (1 - \beta) \cdot v^U \mid Z = 1].$$

The smaller that  $\beta$  is, the more the school places value on the unmeasurable skill  $v^U$ .

We now show that some degree of randomization in the admission policy may be optimal for the ranking designer if they value the unmeasurable skill sufficiently.

**Proposition B.2** (The school's weighted private utility is maximized by some randomization). *Consider the class of two-level policies. For any  $c \in (0, 1 - \rho)$ , there exist some  $\beta \in (0, 1)$  such that the school's utility  $\mathcal{U}_\beta^{\text{pri}}$  is maximized at that value of  $c$ .*

The above result suggests that in the absence of the measurability of one of the skills, the non-randomized admissions policy may have added externality of promoting too much competition in the measurable skill, at



the expense of the unmeasurable one. Our observation is in the same spirit as the finding by [Holmstrom and Milgrom \(1991\)](#) that excessive performance-based incentives can lead strategic agents to focus only on dimensions of achievement that can be effectively measured. A difference between our setting and the aforementioned work is that the excess focus on the measurable skill is being driven by competition between applicants, not (directly) the design  $\lambda$  of the school. However, the school can counter-act this competitive pressure and raise both its welfare and that of the applicants by adding randomization.

**Discussion** In this section, we have only begun to explore strategic ranking in the multi-dimensional setting. We have outlined two promising directions of inquiry: *multi-dimensional competition* and *the lack of measurability*. For future work in the former direction, it would be natural to consider non-linear effort transfer functions, as well as the ramifications for disparate impact when the different environment factors are taken into account. In this setting, we have also highlighted  $\alpha$  as an design choice—it determines how applicants tradeoff effort between the skill dimensions. Different stakeholders, such as the school and the society, may value the skills differently, due to, for example, long-term v.s. short-term considerations, and therefore have different preferences over  $\alpha$ .

The direction regarding measurability is directly motivated by the classic contract design work of [Holmstrom and Milgrom \(1991\)](#), who show that when some work tasks are less measurable than others, it may be optimal to lessen incentives on the measurable tasks (e.g., by adopting a fixed wage) so as to not crowd out effort in the less measurable ones. In our multidimensional setting, the college faces similar tradeoffs, with the additional challenge that it now must rank multiple applicants. For future work, it would be interesting to think about partial measurability and the effect of ‘weak’ measurements on competition. In both cases, a core challenge is characterizing resulting equilibria and analogous rank preservation results.

## C Proofs for Section 2.2

**Lemma C.1.** *In any equilibrium, tie-breaking is not necessary: ties in post-effort scores lead to ties in post-effort rewards. For any  $\Gamma$ , in any equilibrium the distribution of post effort scores  $v$  is such that, for all  $\omega, \omega'$*

$$v(e(\theta_{pre}(\omega)), \theta_{pre}(\omega)) = v(e(\theta_{pre}(\omega')), \theta_{pre}(\omega')) \triangleq v \implies \lambda(\gamma(\omega, v)) = \lambda(\gamma(\omega', v)).$$

*Proof.* We prove the claim by contradiction. Suppose not, and that

$$v(e(\theta_{pre}(\omega)), \theta_{pre}(\omega)) = v(e(\theta_{pre}(\omega')), \theta_{pre}(\omega')) \triangleq v$$

but  $\lambda(\gamma(\omega, v)) < \lambda(\gamma(\omega', v))$ . Then, for any  $\epsilon > 0$ , we have that  $v(e(\theta_{pre}(\omega)) + \epsilon, \theta_{pre}(\omega)) > v$  (as the effort transfer function  $g$  is strictly increasing) and hence

$$\lambda(\gamma(\omega, v)) < \lambda(\gamma(\omega', v)) \leq \lambda(\gamma(\omega, v(e(\theta_{pre}(\omega)) + \epsilon, \theta_{pre}(\omega)))).$$

Since the function  $p$  is continuous, we have that  $p(e(\theta_{pre}(\omega)) + \epsilon) - p(e(\theta_{pre}(\omega))) \rightarrow 0$  as  $\epsilon \rightarrow 0$ .

Then, for small enough  $\epsilon$ , we have that

$$\lambda(\gamma(\omega, v(e(\theta_{pre}(\omega)) + \epsilon, \theta_{pre}(\omega)))) - p(e(\theta_{pre}(\omega)) + \epsilon) > \lambda(\gamma(\omega, v)) - p(e(\theta_{pre}(\omega)))$$

and thus the effort  $e(\theta_{pre}(\omega))$  is not welfare maximizing for the applicant  $\omega$ , a contradiction for it being the equilibrium effort for  $\omega$ .  $\square$

**Lemma C.2** (Deviations of measure 0). *Fix  $\Gamma$  and  $\lambda$ . Consider strategy set  $\{e(\theta_{pre}(\omega))\}$ , and corresponding CDF  $F$  of post-effort scores. As defined,  $\gamma(\omega, v)$  is uniquely determined by  $F$  and  $\Gamma$ . Now, suppose a measure 0 set  $\{\omega\}$  deviates, leading to strategy set  $\{\tilde{e}(\theta_{pre}(\omega))\}$ , post-effort value distribution  $\tilde{F}$ , and ranking function  $\tilde{\gamma}$ .*

*Then,  $\tilde{F} = F$ , and for all  $\omega$  and  $v$ ,*

$$\lambda(\tilde{\gamma}(\omega, v)) = \lambda(\gamma(\omega, v)).$$

*Proof.* By supposition,  $\{e(\theta_{pre}(\omega))\} = \{\tilde{e}(\theta_{pre}(\omega))\}$  except at a set of measure 0, and so the CDFs of the post-effort scores are equal,  $F = \tilde{F}$ . Now, recall that ranking function  $\gamma(\omega, v)$  is defined as the CDF  $F$  except where there are ties of positive mass (atoms) in the distribution of  $v$ .

For  $v$  such that there is not an atom at  $v$ , the equality follows.

If there is an atom at  $v$ , note that for all  $\omega \in \{\omega : v(e(\theta_{pre}(\omega)), \theta_{pre}(\omega)) = v\}$ , we have

$$\gamma(\omega, v) \in [\limsup_{r \uparrow v} F(r), F(v)] = [\limsup_{r \uparrow v} \tilde{F}(r), \tilde{F}(v)].$$

Finally, note that  $\tilde{\gamma}(\omega, v) \in [\limsup_{r \uparrow v} \tilde{F}(r), \tilde{F}(v)]$ , and in particular the CDF of  $\Gamma(w)$  restricted to the atomic set does not change due to measure 0 changes. The equality follows.  $\square$

Lemma C.2 characterizes the effect of measure 0 deviations. While not needed for our results, it establishes that the behavior model in the equilibrium definition is consistent.

**Proposition 2.1:** *In every equilibrium,  $\lambda(\theta_{post}(\theta_{pre})) = \lambda(\theta_{pre})$ , up to sets of measure 0.*

*Proof.* We will prove that in every equilibrium we have that for all  $\omega, \omega'$ ,  $\lambda(\theta_{pre}) > \lambda(\theta'_{pre}) \implies \lambda(\theta_{post}) \geq \lambda(\theta'_{post})$ .

<sup>10</sup>To see that this implies the claim, recall the definition of  $\lambda$  as a function taking  $K$  possible values  $\ell_0, \dots, \ell_{K-1}$ , and consider  $\omega_0, \dots, \omega_{K-1}$  such that  $\lambda(\theta_{pre}(\omega_k))$  runs through the  $K$  possible values for  $\lambda$  in increasing order. The above inequality then implies that the order is weakly preserved if instead of  $\lambda(\theta_{pre}(\omega_k))$  we consider

<sup>10</sup>Here we are using the shorthand notation  $\theta_{pre} := \theta_{pre}(\omega)$  and  $\theta'_{pre} := \theta_{pre}(\omega')$ , as well as  $\theta_{post} := \theta_{post}(\theta_{pre})$  and  $\theta'_{post} := \theta_{post}(\theta'_{pre})$ .

$\lambda(\theta_{\text{post}}(\theta_{\text{pre}}(\omega_k)))$ . However, combined with the fact that  $\theta_{\text{post}}$  is uniformly distributed by our construction of the tie-breaking function, one easily shows the stronger statement that for almost all choices of  $\omega_0, \dots, \omega_{K-1}$ , the order is strictly preserved, which implies that claim of the proposition.

We will prove the above monotonicity claim by contradiction, and thus assume that there exists an equilibrium with a pair  $\omega, \omega'$  with  $\lambda(\theta_{\text{pre}}) > \lambda(\theta'_{\text{pre}})$  such that they spend effort  $e, e'$ , respectively, to end up with post-effort skills  $v, v'$  and  $\lambda(\theta_{\text{post}}) < \lambda(\theta'_{\text{post}})$ . The first inequality implies that  $f(\theta_{\text{pre}}) > f(\theta'_{\text{pre}})$  and the second implies that  $v < v'$ , via Lemma C.1. The high level proof idea for why this is a contradiction is that, if it is worthwhile for  $\theta'_{\text{pre}}$  to spend effort  $e' > e$  to reach skill  $v' > v$  and thus reward  $\lambda(\theta'_{\text{post}})$ , then it would also be worthwhile for  $\theta_{\text{pre}}$  to spend enough effort to reach skill  $v'$  and reward  $\lambda(\theta'_{\text{post}})$ , due to the convexity of the effort cost function  $p(e)$ .

For convenience, in the proof we overload  $\gamma$  to take in applicant effort as opposed to post-effort value as an argument:

$$\gamma(\omega, e) \triangleq \gamma(\omega, v(\theta_{\text{pre}}(\omega), e)).$$

We also assume for convenience in the proof that for the relevant agents, the given effort levels are above the minimum effort level  $e_0$ . A near identical proof follows otherwise.

Then, from the definition of an equilibrium, we have:

$$e \in \underset{d}{\operatorname{argmax}} [\lambda(\gamma(\omega, d)) - p(d)] \quad \text{for } \theta_{\text{pre}} \quad (2)$$

$$\begin{aligned} \theta_{\text{post}} &= \gamma(\omega, e) \\ e' &\in \underset{d}{\operatorname{argmax}} [\lambda(\gamma(\omega', d)) - p(d)] \quad \text{for } \theta'_{\text{pre}} \quad (3) \end{aligned}$$

$$\theta'_{\text{post}} = \gamma(\omega', e')$$

Holding all other effort levels fixed, let  $\tilde{e}$  be the effort that  $\theta_{\text{pre}}$  would have needed to reach skill  $v'$  (by the fact that the effort transfer function is continuous and strictly monotone,  $\tilde{e}$  is uniquely determined by  $v'$ , with  $e < \tilde{e} < e'$ ) and thus rank reward  $\lambda(\theta'_{\text{post}})$  (due to Lemma C.1 and the definition of  $\gamma$ , by deviating to reach a higher score  $v'$ , applicant  $\omega$  receives the same rank reward as  $\omega'$  does in the equilibrium with score  $v'$ ). Similarly let  $\tilde{e}'$ ,  $e < \tilde{e}' < e'$  be the effort that  $\theta'_{\text{pre}}$  would have needed to reach score  $v$  and thus rank reward  $\lambda(\theta_{\text{post}})$ . In other words,  $\lambda(\gamma(\omega, e)) = \lambda(\gamma(\omega', \tilde{e}'))$ , and  $\lambda(\gamma(\omega', e')) = \lambda(\gamma(\omega, \tilde{e}))$ . From Equations (2) and (3), we have:

$$\begin{aligned} \lambda(\gamma(\omega, e)) - p(e) &\geq \lambda(\gamma(\omega, \tilde{e})) - p(\tilde{e}) && \text{Eq. (2)} \\ \text{and } \lambda(\gamma(\omega', e')) - p(e') &\geq \lambda(\gamma(\omega', \tilde{e}')) - p(\tilde{e}') && \text{Eq. (3)} \end{aligned}$$

Applying Lemma C.3, we get  $p(e') - p(\tilde{e}') \leq p(\tilde{e}) - p(e)$ . In addition, Lemma C.4 implies  $e' - \tilde{e}' > \tilde{e} - e$ .

Thus, we have all of the following:

$$\begin{aligned} e' - \tilde{e}' &> \tilde{e} - e \\ p(e') - p(\tilde{e}') &\leq p(\tilde{e}) - p(e) \\ \tilde{e}' &> e \end{aligned}$$

However, they together contradict the assumption that  $p(x)$  is convex and strictly increasing for  $x \geq e_0$ .  $\square$

**Lemma C.3.** *Suppose there are applicants  $\omega, \omega'$  with  $\theta_{\text{pre}}(\omega') < \theta_{\text{pre}}(\omega)$  and effort levels,  $e < \{\tilde{e}', \tilde{e}\} < e'$  such that  $\lambda(\gamma(\omega, e)) = \lambda(\gamma(\omega', \tilde{e}'))$ , and  $\lambda(\gamma(\omega', e')) = \lambda(\gamma(\omega, \tilde{e}))$ . Then the following inequalities*

$$\lambda(\gamma(\omega, e)) - p(e) \geq \lambda(\gamma(\omega, \tilde{e})) - p(\tilde{e}) \quad (4)$$

$$\lambda(\gamma(\omega', e')) - p(e') \geq \lambda(\gamma(\omega', \tilde{e}')) - p(\tilde{e}') \quad (5)$$

imply  $p(e') - p(\tilde{e}') \leq p(\tilde{e}) - p(e)$ .

*Proof.* From Equations (4) and (5), we have

$$\begin{aligned}
 & \lambda(\gamma(\omega, e)) - p(e) \geq \lambda(\gamma(\omega, \tilde{e})) - p(\tilde{e}) \\
 \iff & \lambda(\gamma(\omega, \tilde{e})) - \lambda(\gamma(\omega, e)) \leq p(\tilde{e}) - p(e) && \text{re-arrange} \\
 \\
 \text{and} & \quad \lambda(\gamma(\omega', e')) - p(e') \geq \lambda(\gamma(\omega', \tilde{e}')) - p(\tilde{e}') \\
 \iff & \lambda(\gamma(\omega, \tilde{e})) - p(e') \geq \lambda(\gamma(\omega, e)) - p(\tilde{e}') && \text{defn of } \tilde{e}, \tilde{e}' \\
 \iff & \lambda(\gamma(\omega, \tilde{e})) - \lambda(\gamma(\omega, e)) \geq p(e') - p(\tilde{e}') && \text{re-arrange}
 \end{aligned}$$

Since  $e' > \{\tilde{e}', \tilde{e}\} > e$  and  $e \geq e_0$  in any equilibrium, we have that  $p(e') > \{p(\tilde{e}'), p(\tilde{e})\} > p(e)$ . Thus:

$$0 < p(e') - p(\tilde{e}') \leq \lambda(\gamma(\omega, \tilde{e})) - \lambda(\gamma(\omega, e)) \leq p(\tilde{e}) - p(e)$$

□

**Lemma C.4.** *Suppose there are applicants  $\omega, \omega'$  with  $\theta_{pre}(\omega') < \theta_{pre}(\omega)$  and effort levels,  $e < \{\tilde{e}', \tilde{e}\} < e'$  such that  $v = v(e, \theta_{pre}(\omega)) = v(\tilde{e}', \theta_{pre}(\omega'))$  and  $v' = v(e', \theta_{pre}(\omega')) = v(\tilde{e}, \theta_{pre}(\omega))$ . Then we have that  $e' - \tilde{e}' > \tilde{e} - e$ .*

*Proof.* By assumption, we have  $g(e) = \frac{v}{f(\theta_{pre})}$ ,  $g(e') = \frac{v'}{f(\theta'_{pre})}$ ,  $g(\tilde{e}) = \frac{v'}{f(\theta_{pre})}$ , and  $g(\tilde{e}') = \frac{v}{f(\theta'_{pre})}$ . Thus:

$$\begin{aligned}
 g(\tilde{e}) - g(e) &= \frac{v' - v}{f(\theta_{pre})} \\
 &< \frac{v' - v}{f(\theta'_{pre})} && f(\theta'_{pre}) < f(\theta_{pre}) \\
 &= g(e') - g(\tilde{e}').
 \end{aligned}$$

Since  $g$  is concave and increasing, and  $\tilde{e}' > e$ , we have that

$$g(e') - g(\tilde{e}') > g(\tilde{e}) - g(e) \implies e' - \tilde{e}' > \tilde{e} - e.$$

□

*Proof of Lemma 2.2.* We first prove that the effort levels defined in the theorem lead to  $\lambda(\theta_{pre}) = \lambda(\theta_{post})$ . To this end, we first note that for  $\theta_{pre} \in \psi_k$  the efforts weakly decrease with  $\theta_{pre}$ , while the score,

$$v_k(\theta_{pre}) = \max\{g(\tilde{e}_{k-1})f(c_k), g(e_0)f(\theta_{pre})\}$$

weakly increases with  $\theta_{pre}$ . Thus

$$\begin{aligned}
 \min_{\theta_{pre} \in \psi_k} v_k(\theta_{pre}) &= \max\{g(e_0), g(\tilde{e}_{k-1})\}f(c_k) = g(\tilde{e}_{k-1})f(c_k) \\
 &> g(e_{k-1}(c_k))f(c_k) = \max\{g(e_0)f(c_k), g(\tilde{e}_{k-2})f(c_{k-1})\} \\
 &= v_{k-1}(c_k) = \sup_{\theta_{pre} \in \psi_{k-1}} v_{k-1}(\theta_{pre}).
 \end{aligned}$$

This shows that the scores attained by  $\theta_{pre}$  in  $\psi_k$  are strictly larger than those in  $\psi_{k-1}$ , which in turn implies that  $\lambda(\theta_{pre}) = \lambda(\theta_{post})$ .

Next we note that the effort levels from the theorem have been chosen to guarantee the following: (i) each applicant  $\theta_{pre} \in \psi_k$  exerts the minimal effort that guarantees score  $v_k(c_k)$  or higher, and (ii) for each applicant in  $\psi_{k-1}$ , the effort required to reach score  $v_k(c_k)$  or higher has an additional cost which is equal or larger than the benefit  $\ell_k - \ell_{k-1}$ . Note that these conditions are clearly necessary for the efforts to be an equilibrium – but it is also sufficient for the following two reasons:

1. The construction only directly prevents that an applicant with  $\theta_{\text{pre}} \in \psi_k$  would spend enough effort to match applicants with  $\theta_{\text{pre}} \in \psi_{k+1}$ . However, we also have that they do not strictly benefit from matching applicants with  $\theta_{\text{pre}} \in \psi_j$  for some  $j > k + 1$ . This argument follows from a similar proof by contradiction, as the proof of Proposition 2.1. For  $j = k + 2$ : suppose applicants with  $\theta_{\text{pre}} \in \psi_k$  wish to match applicants with  $\theta_{\text{pre}} \in \psi_{k+2}$ . By construction, applicants with  $\theta_{\text{pre}} \in \psi_{k+1}$  do not prefer to match applicants with  $\theta_{\text{pre}} \in \psi_{k+2}$  over staying in level  $k + 1$ . Also, by construction, applicants with  $\theta_{\text{pre}} \in \psi_k$  do not wish to match applicants with  $\theta_{\text{pre}} \in \psi_{k+1}$  over staying in level  $k$ , but do prefer to move to level  $k + 2$ . Then we may apply Lemmas C.3 and C.4 to obtain a contradiction. Now, iteratively apply the argument for each  $j > k + 2$ .
2. The construction also guarantees that those in band  $\psi_k$  do not profit from reducing their effort, ending up in a lower band  $\psi_j$ , for  $j \leq k - 1$ . First, consider  $j = k - 1$ ; this follows from strict convexity and monotonicity of  $p$  in the relevant region, via an argument similar to that of Proposition 2.1: if an applicant in  $\psi_k$  benefits from moving to a lower band, then an applicant in the lower band strictly benefits to moving to band  $\psi_k$ , which is prevented by construction and the previous argument. More precisely, let  $e'_h$  and  $e'_l$  be the infimum over the effort levels that applicants in band  $k - 1$  need reach the higher and lower bands, respectively, i.e., the effort levels of the applicant with rank  $c_k$ . Let  $e_h, e_l$  similarly be the requisite effort levels for some applicant in  $\psi_k$  with  $\theta_{\text{pre}} > c_k$ . Recall that  $\ell_k \triangleq \lambda(\theta)$  for  $\theta \in \psi_k$ . By construction, we have

$$\ell_k - \ell_{k-1} = p(e'_h) - p(e'_l). \quad (6)$$

As in Proposition 2.1, we further have  $e'_h - e'_l > e_h - e_l$ , and  $e'_l > e_l$ . Since  $p$  is strictly convex and increasing in the region, we have  $p(e_h) - p(e_l) < p(e'_h) - p(e'_l)$ . This fact together with Equation (6) implies  $\ell_k - \ell_{k-1} > p(e_h) - p(e_l)$ , i.e., that the applicant in  $\psi_k$  does not prefer to earn a lower score. Now, for  $j < k - 1$ , iteratively apply the same argument made above in part 1 with Lemmas C.3 and C.4, but with going to lower levels instead of higher levels.

This proves that the effort levels defined in the theorem indeed describe an equilibrium. To complete the proof, we need to prove the converse. To the end, we first observe that by Proposition 2.1, in every equilibrium rank rewards are preserved,  $\lambda(\theta_{\text{pre}}) = \lambda(\theta_{\text{post}})$ . Thus, in equilibrium for every applicant with  $\theta_{\text{pre}} \in \psi_i$ , we can write,

$$\begin{aligned} e_i(\theta_{\text{pre}}) &\in \underset{e}{\operatorname{argmax}} [\lambda(\gamma(\omega, e)) - p(e)] && \text{Defn of equilibrium} && (7) \\ \lambda(\theta_{\text{pre}}) &= \lambda(\theta_{\text{post}}) \triangleq \lambda(\gamma(\omega, e_i(\theta_{\text{pre}}))) && \text{Lemma 2.1} \end{aligned}$$

Consider an applicant with  $\theta_{\text{pre}} \in \psi_i$ , for  $i > 0$ . Fix the effort levels for all applicants with  $\theta'_{\text{pre}} \in \psi_j$ , at some  $e_j(\theta'_{\text{pre}})$ , for  $j < i$ .

Suppose  $e_i(\theta_{\text{pre}}) < g^{-1} \left( \frac{g(\tilde{e}_{i-1}) \cdot f(c_i)}{f(\theta_{\text{pre}})} \right)$ . Then, there exists an applicant with  $\theta'_{\text{pre}} \in \psi_j$ , for  $j = i - 1$ , who could strictly increase their rank and utility by changing their effort level to  $e$  such that  $g^{-1} \left( \frac{g(e_i(\theta_{\text{pre}})) \cdot f(\theta_{\text{pre}})}{f(\theta'_{\text{pre}})} \right) < e < \tilde{e}_{i-1}$ :

$$\begin{aligned} p(e) &= \ell_i - p(e) && g^{-1} \left( \frac{e_i(\theta_{\text{pre}}) \cdot f(\theta_{\text{pre}})}{f(\theta'_{\text{pre}})} \right) < e \\ &> \ell_{i-1} - p(e_{i-1}(\theta'_{\text{pre}})) && e < \tilde{e}_{i-1} \end{aligned}$$

Thus, by contradiction, we have  $e_i(\theta_{\text{pre}}) \geq g^{-1} \left( \frac{g(\tilde{e}_{i-1}) \cdot f(c_i)}{f(\theta_{\text{pre}})} \right)$ .

Now, consider an equilibrium set of effort levels  $\{e_i(\theta_{\text{pre}})\}$ , such that for all  $i$  we have  $e_i(\theta_{\text{pre}}) \geq g^{-1} \left( \frac{g(\tilde{e}_{i-1}) \cdot f(c_i)}{f(\theta_{\text{pre}})} \right)$ .

We note that it follows that  $e_i(\theta_{\text{pre}}) = \max \left( g^{-1} \left( \frac{g(\tilde{e}_{i-1}) \cdot f(c_i)}{f(\theta_{\text{pre}})} \right), \operatorname{argmin}_e p(e) \right)$ : for each applicant  $\theta_{\text{pre}}$ , applicants in higher tiers exert enough effort such that  $\theta_{\text{pre}}$  does not wish to exert enough effort to overtake them, and tying them does not strictly increase their welfare. Thus, these applicants minimize their effort cost while still exerting enough effort to prevent lower tier applicants from overtaking them. Note that it is not an issue that  $\theta_{\text{pre}}$  is *indifferent* in overtaking higher tiers. In any equilibrium, only a measure 0 set of such lower tier applicants will exert higher effort to tie the higher tier. Otherwise, tie-breaking will be necessary, in violation of Lemma C.1.

We finish the proof by showing there exists a unique effort set  $\{e_i\}$  following the above. For applicants in  $\psi_0$ ,  $\theta_{\text{pre}} = \theta_{\text{post}} \in \psi_0$ , and so  $e_0(\theta_{\text{pre}}) = \operatorname{argmin}_e(p(e))$  in every equilibrium. Applicants do not need to spend any effort to stay at rank 0, and in equilibrium they will not be incentivized to exert enough effort to increase their admissions probability. Recursively, efforts in higher tiers are fixed given effort in lower tiers, and so we have constructed the unique equilibrium.  $\square$

**Proof for Corollary A.1** Assume that  $g$  and  $p$  are differentiable. If  $\ell_k$  increases and  $\ell_j$  decreases for some  $j > k$  (holding all other parameters fixed), then  $e_i(\theta)$  for all  $i < k - 1$  are unaffected,  $e_k$  is weakly increasing, and the efforts  $e_{k+1}, \dots, e_j$  are weakly decreasing.

*Proof.* By the inductive definitions of efforts, those for applicants  $\theta_{\text{pre}} \in \psi_i$  for  $i < k$  do not change; to calculate the change of the effort for applicants in  $\psi_k$ , we differentiate the equation defining  $\tilde{e}_{k-1}$ , giving  $p'(\tilde{e}_{k-1}) \frac{d}{d\ell_k} \tilde{e}_{k-1} = 1$ , and hence

$$\frac{d\tilde{e}_{k-1}}{d\ell_k} = \frac{1}{p'(\tilde{e}_{k-1})} > 0.$$

This implies that  $\tilde{e}_{k-1}$  and hence  $e_k$  is increasing in  $\ell_k$ . Next we differentiate the defining equation for  $\tilde{e}_k$  with respect to  $\ell_k$ , yielding

$$p'(\tilde{e}_k) \frac{d\tilde{e}_k}{d\ell_k} = p'(e_k(c_{k+1})) \frac{de_k(c_{k+1})}{d\ell_k} - 1$$

Using that  $g(e_k(c_{k+1})) = \max \left\{ g(e_0), g(\tilde{e}_{k-1}) \frac{f(c_k)}{f(c_{k+1})} \right\}$  we know that if  $\frac{de_k(c_{k+1})}{d\ell_k} \neq 0$ , then

$$g(e_k(c_{k+1})) = g(\tilde{e}_{k-1}) \frac{f(c_k)}{f(c_{k+1})} < g(\tilde{e}_{k-1}) \quad \text{and} \quad \frac{de_k(c_{k+1})}{d\ell_k} = \frac{f(c_k)}{f(c_{k+1})} \frac{g'(\tilde{e}_{k-1})}{g'(e_k(c_{k+1}))} \frac{d\tilde{e}_{k-1}}{d\ell_k}$$

where the second equality follows from the first and the chain rule. The first bound implies that  $e_k(c_{k+1}) < \tilde{e}_{k-1}$  so by the concavity of  $g$ ,  $g'(\tilde{e}_{k-1}) \leq g'(e_k(c_{k+1}))$  and by the convexity of  $p$ ,  $p'(\tilde{e}_{k-1}) \geq p'(e_k(c_{k+1}))$ ; as a consequence

$$\frac{de_k(c_{k+1})}{d\ell_k} < \frac{d\tilde{e}_{k-1}}{d\ell_k} = \frac{1}{p'(\tilde{e}_{k-1})} \leq \frac{1}{p'(e_k(c_{k+1}))},$$

and hence

$$\frac{d}{d\ell_k} \tilde{e}_k = \frac{1}{p'(\tilde{e}_k)} \left( p'(e_k(c_{k+1})) \frac{de_k(c_{k+1})}{d\ell_k} - 1 \right) < 0.$$

This shows that  $e_{k+1}(\theta)$  is weakly decreasing in  $\ell_k$ . By the inductive definitions, the efforts  $e_{k+2}, \dots, e_{j-1}$  are non-increasing, while  $\tilde{e}_{j-1}$  is strictly decreasing, due to the change of  $\ell_j$ . This in turn implies that  $e_j$  is weakly decreasing in  $\psi_j$ .  $\square$

## D Proofs for Section 3

### Proof of Proposition 3.1.

*Proof.* We prove this result under more general conditions where  $e_0 \geq 0$ .

For a  $K$ -level policy, we can write the applicant welfare as follows:

$$\begin{aligned} \mathcal{W} &= \sum_{i=0}^{K-1} \mathbb{E} [\ell_i - p(e_i(\theta_{\text{pre}})) \mid \theta_{\text{pre}} \in \psi_i] \\ &= \rho - \sum_{i=0}^{K-1} \mathbb{E} [p(e_i(\theta_{\text{pre}})) \mid \theta_{\text{pre}} \in \psi_i] \\ &\leq \rho, \end{aligned}$$

since the term  $\mathbb{E}[p(e_i(\theta_{\text{pre}})) \mid \theta_{\text{pre}} \in \psi_i] \geq 0$  for any  $i$  such that  $0 \leq i \leq K - 1$ .

The maximum  $\mathcal{W}$  is uniquely attained by a two-level  $\lambda$  where  $c_1 = 0$ ,  $e_i(\theta_{\text{pre}}) = e_0$ , and so  $p(e_i(\theta_{\text{pre}})) = 0$  for all  $\theta_{\text{pre}}$ .

In the two-level policy class:

$$\begin{aligned} \mathcal{W} &= (1 - c) \cdot \mathbb{E} \left[ \ell_1 - p(\max(g^{-1}(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}})}), e_0)) \mid \theta_{\text{pre}} > c \right] \\ &= \rho - \int_c^1 p(\max(g^{-1}(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}})}), e_0)) d\theta_{\text{pre}}. \end{aligned}$$

Taking the derivative of the above display:

$$\frac{\partial \mathcal{W}}{\partial c} = -p(\max(\tilde{e}_0, e_0)) - \int_c^1 \frac{\partial}{\partial c} p(\max(g^{-1}(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}})}), e_0)) d\theta_{\text{pre}}$$

which is non-positive since  $\frac{\partial}{\partial c} p(\max(g^{-1}(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}})}), e_0)) \geq 0$ . □

### Proof of Proposition 3.2.

*Proof.* Note that for a two-level policy with cutpoint  $c$ , we have  $\tilde{e}_0 = p^{-1}(\ell_1) = p^{-1}(\frac{\rho}{1-c})$ . We have, since  $g(e_0) = 0$ ,

$$\begin{aligned} \mathbb{E}[v^M \mid Z = 1] &= g(\tilde{e}_0)f(c) \\ &= g(p^{-1}(\frac{\rho}{1-c}))f(c) \end{aligned}$$

which is increasing in  $c$ . Thus  $\mathcal{U}^{\text{pri}}$  is maximized by choosing  $c$  to be as large as possible while filling the school's capacity, that is,  $c = 1 - \rho$ , which corresponds to  $\ell_1 = 1$ . □

### Proof of Proposition 3.3.

*Proof.* Consider the following setting: The effort transfer function  $g$  and effort cost function  $p$  are such that  $g(p^{-1}(x)) = x$  and  $g(e_0) = 0$ . We will show the following claim:

“For any three-level  $\lambda$  with levels  $(\ell_0, \ell_1, \ell_2) = (0, x, 1)$  and cutpoints  $0 < c_1 < c_2 < c_3 = 1$ , there exists a skill distribution with quantile function  $f$  such that  $\mathcal{U}^{\text{pri}}$  is higher under  $\lambda$  than under the two-level policy corresponding to non-randomization.”

To show that  $\mathcal{U}^{\text{pri}}$  is higher under the three-level policy, we show the following inequality:

$$\begin{aligned} &g(\tilde{e}_0) \cdot f(c_1) \cdot (c_2 - c_1) \cdot x + g(p^{-1}(1 - x + pg^{-1}(\frac{g(\tilde{e}_0) \cdot f(c_1)}{f(c_2)}))) \cdot f(c_2) \cdot (1 - c_2) \\ &> g(p^{-1}(1)) \cdot f(1 - \rho) \cdot \rho, \end{aligned} \tag{8}$$

where  $\tilde{e}_0 = p^{-1}(x)$ . The LHS is nothing but  $\mathcal{U}^{\text{pri}}$  under a three-level policy, and the RHS,  $\mathcal{U}^{\text{pri}}$  under the two-level policy with cutpoint  $1 - \rho$ .

Using  $g(p^{-1}(x)) = x$ , we can rewrite (8) as follows:

$$x \cdot f(c_1) \cdot (c_2 - c_1) \cdot x + ((1 - x)f(c_2) + x \cdot f(c_1)) \cdot (1 - c_2) > f(1 - \rho) \cdot \rho. \quad (9)$$

The above display is true whenever  $f(c_2)$  is large enough, that is

$$f(c_2) > \frac{f(1 - \rho) \cdot \rho - x^2 \cdot f(c_1) \cdot (c_2 - c_1) - x \cdot f(c_1) \cdot (1 - c_2)}{(1 - x)(1 - c_2)}.$$

□

### Proof of Proposition 3.4

*Proof.* Given that  $g(e_0) = 0$ , we can rewrite the societal utility as

$$\begin{aligned} \mathcal{U}^{\text{soc}} &= \mathbb{E}[v^M] \\ &= (1 - c) \cdot g(\tilde{e}_0) \cdot f(c), \end{aligned}$$

which attains the value 0 at  $c = 0$  and  $c = 1$ . Therefore it has an interior maximizer  $c^* \in (0, 1 - \rho)$ , if and only if

$$\frac{\partial}{\partial c}(1 - c) \cdot g(p^{-1}(\frac{\rho}{1 - c})) \cdot f(c) < 0, \text{ at } c = 1 - \rho$$

It suffices to check that the above inequality is true for

$$f(x) = x, \quad g(x) = \sqrt{x}, \quad p(x) = x^2, \quad \rho < 5/9.$$

□

## E Comparison table for Section 4

Latent skill rank $\theta_{\text{true}}$	“Low” region $0 \leq \theta_{\text{true}} < \text{thres}_A(c)$	“Middle” region $\text{thres}_A(c) \leq \theta_{\text{true}} < \text{thres}_B(c)$	“High” region $\text{thres}_B(c) \leq \theta_{\text{true}} \leq 1$
Welfare	$\mathcal{W}^A(\theta_{\text{true}}) = \mathcal{W}^B(\theta_{\text{true}}) = 0$	$\mathcal{W}^A(\theta_{\text{true}}) \geq \mathcal{W}^B(\theta_{\text{true}}) = 0$	$\mathcal{W}^A(\theta_{\text{true}}) > \mathcal{W}^B(\theta_{\text{true}}) \geq 0$
Admission probability $\lambda(\theta_{\text{post}})$	$\lambda^A(\theta_{\text{true}}) = \lambda^B(\theta_{\text{true}}) = 0$	$\lambda^A(\theta_{\text{true}}) = \frac{\rho}{1 - c} > \lambda^B(\theta_{\text{true}}) = 0$	$\lambda^A(\theta_{\text{true}}) = \lambda^B(\theta_{\text{true}}) = \frac{\rho}{1 - c}$
Effort	$e^A(\theta_{\text{true}}) = e^B(\theta_{\text{true}}) = e_0$	$e^A(\theta_{\text{true}}) > e^B(\theta_{\text{true}}) = e_0$	$e^B(\theta_{\text{true}}) > e^A(\theta_{\text{true}}) > e_0$

Table 1: Comparison of welfare, admission probability and effort between group A and B for every  $\theta_{\text{true}}$ , under a two-level policy parametrized by  $c$ . In the table we have used the shorthand notation  $\lambda^A(\theta_{\text{true}})$  for the admission probability  $\lambda(\theta_{\text{post}}(\theta_{\text{true}}, \psi_A))$  and  $e^A(\theta_{\text{true}})$  for the effort  $e(\theta_{\text{true}}, \psi_A)$ , and the respective notation for group B.

## F Proofs for Section 4

### Proof of Proposition 4.1

*Proof.* In order to transform the current setting into the setting of Proposition 2.1, we absorb the environment factor  $\psi$  into the applicant’s latent skill level  $f(\theta_{\text{true}})$  and compute the applicant’s rank in  $f(\theta_{\text{true}}) \cdot \psi$ . One can check that  $f_{\text{mix}}^{-1}(\cdot)$  as defined is indeed the appropriate CDF for the distribution of  $f(\theta_{\text{true}}) \cdot \psi$  in the overall applicant population. We can now write  $v = g(e) \cdot f_{\text{mix}}(\theta_{\text{pre}})$  and apply Proposition 2.1 to  $\theta_{\text{pre}}$ . □

### Proof of Proposition 4.2



*Proof.* The first part of the result follows from Proposition 4.1. Now to see that the statement about the welfare gap is true, we note the following: (1) For  $0 \leq \theta_{\text{true}} < \text{thres}_A(c)$ , both the group A and group B applicant have zero welfare because they have zero probability of admission and incur zero cost of effort; (2) For  $\text{thres}_A(c) \leq \theta_{\text{true}} < \text{thres}_B(c)$ , the group B has zero welfare (since they have zero probability of admission) whereas the group A with the same  $\theta_{\text{true}}$  has positive probability of admission and non-negative welfare; (3) For  $\text{thres}_B(c) \leq \theta_{\text{true}} \leq 1$ , the group A and group B applicant with same  $\theta_{\text{true}}$  both have the same probability of admission but the group A applicant exerts less effort. This is because a applicant with environment-scaled rank  $\theta_{\text{pre}}$  exerts effort

$$e = \begin{cases} g^{-1}\left(\frac{g(\tilde{e}_0) \cdot f_{\text{mix}}(c)}{f_{\text{mix}}(\theta_{\text{pre}})}\right) & \text{if } \theta_{\text{pre}} > c \\ e_0 & \text{o.w.} \end{cases} \quad (10)$$

For the same  $\theta_{\text{true}}$ , the group A applicant has higher  $\theta_{\text{pre}}$ , since

$$f_{\text{mix}}^{-1}(f(\theta_{\text{true}}) \cdot \psi_A) > f_{\text{mix}}^{-1}(f(\theta_{\text{true}}) \cdot \psi_B).$$

Since  $e$  is decreasing in  $\theta_{\text{pre}}$ , the group A applicant exerts less effort than the group B applicant.  $\square$

### Proof of Proposition 4.3

*Proof.* We have that

$$\frac{\partial \mathcal{G}(\theta_{\text{true}})}{\partial c} = \frac{\partial}{\partial c} \mathcal{W}^A(\theta_{\text{true}}) - \frac{\partial}{\partial c} \mathcal{W}^B(\theta_{\text{true}})$$

Since  $c < (f_{\text{mix}})^{-1}(f(\theta_{\text{true}}) \cdot \psi_B)$ , we have that  $\theta_{\text{true}} > \text{thres}_B(c)$ . Thus we may compute  $\mathcal{W}^A(\theta_{\text{true}})$  and  $\frac{\partial}{\partial c} \mathcal{W}^B(\theta_{\text{true}})$  as follows:

$$\begin{aligned} \mathcal{W}^A(\theta_{\text{true}}) &= \frac{\rho}{1-c} - p\left(g^{-1}\left(\frac{g(\tilde{e}_0) f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_A}\right)\right), \\ \mathcal{W}^B(\theta_{\text{true}}) &= \frac{\rho}{1-c} - p\left(g^{-1}\left(\frac{g(\tilde{e}_0) f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_B}\right)\right), \end{aligned}$$

where as before  $\tilde{e}_0 = p^{-1}\left(\frac{\rho}{1-c}\right)$ . Then we have:

$$\begin{aligned} \frac{\partial}{\partial c} \mathcal{W}^A(\theta_{\text{true}}) &= \frac{\rho}{(1-c)^2} \\ &\quad - p'\left(g^{-1}\left(\frac{g(\tilde{e}_0) f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_A}\right)\right) \cdot (g^{-1})'\left(\frac{g(\tilde{e}_0) f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_A}\right) \cdot \frac{g(\tilde{e}_0)(f_{\text{mix}})'(c) + f_{\text{mix}}(c)g'(\tilde{e}_0) \frac{\partial \tilde{e}_0}{\partial c}}{f(\theta_{\text{true}}) \cdot \psi_A} \end{aligned}$$

We claim that  $\frac{\partial}{\partial c} \mathcal{W}^A(\theta_{\text{true}}) > \frac{\partial}{\partial c} \mathcal{W}^B(\theta_{\text{true}})$ . We have the following facts:

$$\frac{1}{\psi_A} < \frac{1}{\psi_B} \quad (11)$$

$$p'(x) \leq p'(\tilde{x}) \quad \forall x < \tilde{x} \quad (12)$$

$$g^{-1}(x) < g^{-1}(\tilde{x}) \quad \forall x < \tilde{x} \quad (13)$$

$$(g^{-1})'(x) \leq (g^{-1})'(\tilde{x}) \quad \forall x < \tilde{x} \quad (14)$$

$$g(x), g'(x), f_{\text{mix}}(x), f(x) > 0 \quad \forall x > 0 \quad (15)$$

$$(f_{\text{mix}})'(c) > 0 \quad (16)$$

$$\frac{\partial \tilde{e}_0}{\partial c} > 0 \quad (17)$$

$$(18)$$

Equations (11), (12), (13) and (15) imply

$$p'(g^{-1}(\frac{g(\tilde{e}_0)f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_{\text{A}}})) \leq p'(g^{-1}(\frac{g(\tilde{e}_0)f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_{\text{B}}})) \quad (19)$$

Equations (11), (14) and (15) imply

$$(g^{-1})'(\frac{g(\tilde{e}_0)f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_{\text{A}}}) \leq (g^{-1})'(\frac{g(\tilde{e}_0)f_{\text{mix}}(c)}{f(\theta_{\text{true}}) \cdot \psi_{\text{B}}}) \quad (20)$$

Equations (11), (15), (17) and (16) imply

$$\frac{g(\tilde{e}_0)(f_{\text{mix}})'(c) + f_{\text{mix}}(c)g'(\tilde{e}_0)\frac{\partial \tilde{e}_0}{\partial c}}{f(\theta_{\text{true}}) \cdot \psi_{\text{A}}} < \frac{g(\tilde{e}_0)(f_{\text{mix}})'(c) + f_{\text{mix}}(c)g'(\tilde{e}_0)\frac{\partial \tilde{e}_0}{\partial c}}{f(\theta_{\text{true}}) \cdot \psi_{\text{B}}} \quad (21)$$

Intermediate equations (19), (20) and (21) together imply that  $\frac{\partial}{\partial c} \mathcal{W}^{\text{A}}(\theta_{\text{true}}) > \frac{\partial}{\partial c} \mathcal{W}^{\text{B}}(\theta_{\text{true}})$ . This gives  $\frac{\partial \mathcal{G}(\theta_{\text{true}})}{\partial c} > 0$  as desired.  $\square$

### Proof of Proposition 4.4

*Proof.* For a two-level policy with cutpoint  $c > 0$ , we have, by definition:

$$\mathcal{A} = \frac{\rho}{1-c} (1 - \text{thres}_{\text{B}}(c)).$$

Since  $\text{thres}_{\text{B}}(c) > c$  for any  $c \in (0, 1 - \rho]$ , we have that  $\mathcal{A} < \rho$ . On the other hand,  $\mathcal{A} = 0$  for a pure randomization policy.

Taking derivative of  $\mathcal{A}$  with respect to  $c$ ,

$$\begin{aligned} \frac{\partial \mathcal{A}}{\partial c} &= \frac{\rho}{1-c} \left( -\frac{\partial}{\partial c} \text{thres}_{\text{B}}(c) \right) + \frac{\rho}{(1-c)^2} (1 - \text{thres}_{\text{B}}(c)) \\ &= \frac{\rho}{1-c} \left( -\frac{\partial}{\partial c} \text{thres}_{\text{B}}(c) + \frac{1 - \text{thres}_{\text{B}}(c)}{1-c} \right). \end{aligned}$$

We want to show that the above derivative is non-positive under the assumption that  $f^{-1}$  is convex. Since  $\frac{1 - \text{thres}_{\text{B}}(c)}{1-c} \leq 1$  (because  $\text{thres}_{\text{B}}(c) \leq c$  for any  $c \in [0, 1 - \rho]$ ), it suffices to show that  $\frac{\partial}{\partial c} \text{thres}_{\text{B}}(c) \geq 1$ . Indeed, we may compute

$$\begin{aligned} \frac{\partial}{\partial c} \text{thres}_{\text{B}}(c) &= \frac{(f^{-1})' \left( \frac{f_{\text{mix}}(c)}{\psi_{\text{B}}} \right)}{\psi_{\text{B}}} \cdot \left( \frac{1}{2} \cdot \frac{(f^{-1})' \left( \frac{f_{\text{mix}}(c)}{\psi_{\text{A}}} \right)}{\psi_{\text{A}}} + \frac{1}{2} \cdot \frac{(f^{-1})' \left( \frac{f_{\text{mix}}(c)}{\psi_{\text{B}}} \right)}{\psi_{\text{B}}} \right)^{-1} \\ &\geq 1, \end{aligned}$$

since  $\frac{(f^{-1})' \left( \frac{f_{\text{mix}}(c)}{\psi_{\text{B}}} \right)}{\psi_{\text{B}}} \geq \frac{(f^{-1})' \left( \frac{f_{\text{mix}}(c)}{\psi_{\text{A}}} \right)}{\psi_{\text{A}}}$  by the convexity of  $f^{-1}$  (which implies  $(f^{-1})'$  is non-decreasing) and the fact that  $\psi_{\text{A}} > \psi_{\text{B}}$ .  $\square$

## G Proofs for Section B

### Proof of Proposition B.1

*Proof.* The students are sorted according to their weighted scores:

$$\begin{aligned} v^{\alpha} &= \sum_{i=1}^m \alpha_i g(e_i) \cdot f_i(\theta_{\text{pre}}^i) \\ &= h \sum_{i=1}^m e_i \cdot (\alpha_i f_i(\theta_{\text{pre}}^i)). \end{aligned}$$

Consider a single applicant. Denote  $i^* := \operatorname{argmax}_i \alpha_i f_i(\theta_{\text{pre}}^i)$ . Since  $f_i$ 's are strictly monotone and  $\theta_{\text{pre}}^i$  are independently distributed,  $i^*$  is almost everywhere unique. For any  $e > 0$ , let  $e_{i^*}^* = e$  and  $e_i^* = 0$  for each  $i \neq i^*$ . Then, holding the effort levels of all other applicants fixed, we have

$$W(\{e_i^*\}_{i=1}^m, \lambda(\theta_{\text{post}}^\alpha)) \geq W(\{e_i\}_{i=1}^m, \lambda(\theta_{\text{post}}^\alpha)) \quad \forall \{e_i\}_{i=1}^m \text{ s.t. } \sum_{i=1}^m e_i = e.$$

In other words, for any fixed  $\sum_{i=1}^m e_i = e$ , the applicant maximizes their welfare by putting all effort into the skill  $i^*$  with the highest  $\alpha^i \theta_{\text{pre}}^i$ . Therefore, it suffices for the applicant to maximize their individual welfare over the total effort level  $e$ . At equilibrium, the weighted score  $v^\alpha$  satisfies

$$v^\alpha = h \cdot e \cdot v_{\text{pre}}^\alpha.$$

Since  $\lambda$  is non-decreasing in the weighted score  $v^\alpha$ , we may now retrace the proof of Proposition 2.1 to argue the following: if a student with  $v_{\text{pre}}^\alpha$  finds it optimal to achieve weighted score  $v^\alpha$ , then each student with  $\overline{v_{\text{pre}}^\alpha}$  finds it optimal to reach weighted score  $\overline{v^\alpha} \geq v^\alpha$ , due to the convexity of the effort cost function  $p(e)$ .  $\square$

### Proof of Proposition B.2

*Proof.* We prove the result for  $g(e_0) = 0$  to simplify the presentation. The generalization to  $g(e_0)$  can be accomplished by some additional book-keeping without modifying the core idea.

Note that in this setting we have  $\tilde{e}_0 = p^{-1}(\ell_1) = p^{-1}\left(\frac{\rho}{1-c}\right)$ .

For the measurable skill:

$$\mathbb{E}[v^M \mid Z = 1] = g(\tilde{e}_0) \cdot f(c).$$

For the unmeasurable skill:

$$\begin{aligned} \mathbb{E}[v^U \mid Z = 1] &= \mathbb{E}[g(B - e_1(\theta_{\text{pre}}^M)) \cdot f(\theta_{\text{pre}}^M) \mid Z = 1] \\ &= \mathbb{E}\left[g\left(B - g^{-1}\left(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}}^M)}\right)\right) \cdot f(\theta_{\text{pre}}^U) \mid \theta_{\text{pre}}^M > c\right] \\ &= \mathbb{E}[\theta_{\text{pre}}^U] \cdot \int_c^1 g\left(B - g^{-1}\left(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}}^M)}\right)\right) d\theta_{\text{pre}}^M \quad (\text{independence of } \theta_{\text{pre}}^U \text{ and } \theta_{\text{pre}}^M) \end{aligned}$$

We know that  $\frac{\partial}{\partial c} \mathbb{E}[v^M \mid Z = 1] > 0$  because  $\tilde{e}_0$  is increasing in  $c$ .

By the Leibniz integral rule, we have

$$\frac{\partial}{\partial c} \mathbb{E}[v^U \mid Z = 1] = \mathbb{E}[\theta_{\text{pre}}^U] \cdot \left(-g(B - \tilde{e}_0) + \int_c^1 \frac{\partial}{\partial c} g\left(B - g^{-1}\left(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}}^M)}\right)\right) d\theta_{\text{pre}}^M\right) < 0,$$

since  $g\left(B - g^{-1}\left(\frac{g(\tilde{e}_0)f(c)}{f(\theta_{\text{pre}}^M)}\right)\right)$  is decreasing in  $c$ .

For any  $c \in (0, 1 - \rho)$ , set

$$\alpha = \frac{-\frac{\partial}{\partial c} \mathbb{E}[v^U \mid Z = 1]}{\frac{\partial}{\partial c} \mathbb{E}[v^M \mid Z = 1] - \frac{\partial}{\partial c} \mathbb{E}[v^U \mid Z = 1]}.$$

Then we have that

$$\frac{\partial}{\partial c} \mathcal{U} = \alpha \frac{\partial}{\partial c} \mathbb{E}[v^M \mid Z = 1] + (1 - \alpha) \frac{\partial}{\partial c} \mathbb{E}[v^U \mid Z = 1] = 0.$$

That is, the school's utility  $\mathcal{U}$  for the chosen  $\alpha$  is maximized at an interior skill of  $c$ .

$\square$

## References

- R. Abebe and K. Goldner. Mechanism design for social good. *AI Matters*, 4(3):27–34, 2018. 3
- T. Alon, M. Dobson, A. Procaccia, I. Talgam-Cohen, and J. Tucker-Foltz. Multiagent evaluation mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1774–1781, 2020. 2
- H. Aziz. Developments in multi-agent fair allocation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13563–13568, 2020. 3
- Y. Barut and D. Kovenock. The symmetric multiple prize all-pay auction with complete information. *European Journal of Political Economy*, 14(4):627–644, 1998. 2
- Y. Bechavod, K. Ligett, S. Wu, and J. Ziani. Gaming helps! learning from strategic interactions in natural dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 1234–1242. PMLR, 2021. 2
- G. S. Becker. A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846, 1973. 3
- A. L. Bodoh-Creed and B. R. Hickman. College assignment as a large contest. *Journal of Economic Theory*, 175:88–126, 2018. 1, 2, 6, 13
- M. Braverman and S. Garg. The role of randomness and noise in strategic classification. *arXiv preprint arXiv:2005.08377*, 2020. 2, 9, 13
- M. Brückner, C. Kanzow, and T. Scheffer. Static prediction games for adversarial learning problems. *The Journal of Machine Learning Research*, 13(1):2617–2654, 2012. 1, 2
- W. Cai, J. Gaebler, N. Garg, and S. Goel. Fair allocation through selective information acquisition. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES '20*, page 2228, New York, NY, USA, 2020. Association for Computing Machinery. 3
- C. Calsamiglia. Decentralizing equality of opportunity\*. *International Economic Review*, 50(1):273–290, 2009. 3, 7
- A. Chouldechova and A. Roth. The Frontiers of Fairness in Machine Learning. *CoRR*, abs/1810.08810, 2018. 3
- B. L. Connelly, L. Tihanyi, T. R. Crook, and K. A. Gangloff. Tournament theory: Thirty years of contests and competitions. *Journal of Management*, 40(1):16–47, 2014. 2
- L. C. Corchón. The theory of contests: a survey. *Review of economic design*, 11(2):69–100, 2007. 2
- J. Dong, A. Roth, Z. Schutzman, B. Waggoner, and Z. S. Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018. 2
- E. Elkind, A. Ghosh, and P. Goldberg. Contest design with threshold objectives, 2021. 3
- Y. Faenza, S. Gupta, and X. Zhang. Impact of bias on school admissions and targeted interventions. *arXiv preprint arXiv:2004.10846*, 2020. 3
- D. Fang, T. Noe, and P. Strack. Turning up the heat: The discouraging effect of competition in contests. *Journal of Political Economy*, 128(5):1940–1975, 2020. 2
- Q. Fu and Z. Wu. Contests: Theory and topics. In *Oxford Research Encyclopedia of Economics and Finance*. 2019. 2
- N. Garg, H. Li, and F. Monachou. Dropping standardized testing for admissions: Differential variance and access. *arXiv preprint arXiv:2010.04396*, 2020. 3
- P. Gözl, A. Kahng, and A. D. Procaccia. Paradoxes in fair machine learning. *NeurIPS'19*, 2019. 3
- N. Haghtalab, N. Immorlica, B. Lucier, and J. Z. Wang. Maximizing welfare with incentive-aware evaluation mechanisms. *arXiv preprint arXiv:2011.01956*, 2020. 2
- E. A. Hanushek and L. Woessmann. *The high cost of low educational performance: The long-run economic impact of improving PISA outcomes*. ERIC, 2010. 4
- M. Hardt, N. Megiddo, C. Papadimitriou, and M. Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS '16*, pages 111–122, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4057-1. 1, 2
- B. Holmstrom and P. Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.*, 7:24, 1991. 16, 17

- L. Hu and Y. Chen. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference, WWW '18*, pages 1389–1398, Republic and Canton of Geneva, Switzerland, 2018. International World Wide Web Conferences Steering Committee. ISBN 978-1-4503-5639-8. 3
- L. Hu, N. Immorlica, and J. W. Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, pages 259–268, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6125-5. doi: 10.1145/3287560.3287597. 2, 9
- N. Immorlica, K. Ligett, and J. Ziani. Access to population-level signaling as a source of inequality. In *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*, pages 249–258, 2019. 3
- S. Kannan, A. Roth, and J. Ziani. Downstream effects of affirmative action. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 240–248, 2019. 3
- S. Kannan, M. Niu, A. Roth, and R. Vohra. Best vs. all: Equity and accuracy of standardized test score reporting. *arXiv preprint arXiv:2102.07809*, 2021. 3
- J. Kleinberg and M. Raghavan. How Do Classifiers Induce Agents to Invest Effort Strategically? In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC '19*, pages 825–844, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6792-9. 2, 15
- K. A. Konrad. Strategy in contests—an introduction. *WZB-Markets and Politics Working Paper No. SP II, 1*, 2007. 2
- I. Lindenlaub and F. Postel-Vinay. Multidimensional sorting under random search. *Manuscript, University College London*, 2016. 15
- L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pages 3150–3158. PMLR, 2018. 3
- L. T. Liu, A. Wilson, N. Haghtalab, A. T. Kalai, C. Borgs, and J. Chayes. The disparate equilibria of algorithmic decision making when individuals invest rationally. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 381–391, 2020. 3
- Z. Liu and N. Garg. Test-optional policies: Overcoming strategic behavior and informational gaps. *arXiv preprint arXiv:2107.08922*, 2021. 3
- M. Mathioudakis, C. Castillo, G. Barnabo, and S. Celis. Affirmative action policies for top-k candidates selection: with an application to the design of policies for university admissions. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, pages 440–449, 2020. 3
- J. Miller, S. Milli, and M. Hardt. Strategic classification is causal modeling in disguise. In *International Conference on Machine Learning*, pages 6917–6926. PMLR, 2020. 2
- S. Milli, J. Miller, A. D. Dragan, and M. Hardt. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, pages 230–239, New York, NY, USA, 2019. ACM. 2, 9
- H. Mouzannar, M. I. Ohanessian, and N. Srebro. From fair decision making to social equality. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT\* '19*, pages 359–368, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6125-5. 3
- R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. 5, 13
- A. Noriega-Campero, M. A. Bakker, B. Garcia-Bulle, and A. Pentland. Active fairness in algorithmic decision making. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 77–83, 2019. 3
- W. Olszewski and R. Siegel. Large contests. *Econometrica*, 84(2):835–854, 2016. 1, 2, 3, 6, 13
- W. Olszewski and R. Siegel. Bid caps in large contests. *Games and Economic Behavior*, 115:101–112, 2019a. 1, 2, 3
- W. Olszewski and R. Siegel. Pareto improvements in the contest for college admissions. *Unpublished paper, Department of Economics, Northwestern University.[1607]*, 2019b. 1, 2, 3, 13
- J. E. Roemer. *Equality of opportunity*. Harvard University Press, 1998. 3, 7
- E. Rolf, M. Simchowitz, S. Dean, L. T. Liu, D. Bjorkegren, M. Hardt, and J. Blumenstock. Balancing competing objectives with noisy data: Score-based classifiers for welfare-aware machine learning. In *International Conference on Machine Learning*, pages 8158–8168. PMLR, 2020. 3

- D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, Apr 1973. 3
- Y. Shavit, B. Edelman, and B. Axelrod. Causal strategic linear regression. In *International Conference on Machine Learning*, pages 8676–8686. PMLR, 2020. 2
- M. Spence. Job market signaling. In *Uncertainty in economics*, pages 281–306. Elsevier, 1978. 2, 3
- B. Tabibian, V. Gomez, A. De, B. Schölkopf, and M. G. Rodriguez. On the design of consequential ranking algorithms. In *Conference on Uncertainty in Artificial Intelligence*, pages 171–180. PMLR, 2020. 3
- M. Zehlike and C. Castillo. Reducing disparate exposure in ranking: A learning to rank approach. In *Proceedings of The Web Conference 2020*, pages 2849–2855, 2020. 3
- M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. Baeza-Yates. Fa\* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1569–1578, 2017. 3