



HAL
open science

Fairness in Selection Problems with Strategic Candidates

Vitalii Emelianov, Nicolas Gast, Patrick Loiseau

► **To cite this version:**

Vitalii Emelianov, Nicolas Gast, Patrick Loiseau. Fairness in Selection Problems with Strategic Candidates. EC 2022 - 23th ACM Conference on Economics and Computation, Jul 2022, Boulder, Colorado, United States. pp.1-29, 10.1145/3490486.3538287 . hal-03677966

HAL Id: hal-03677966

<https://inria.hal.science/hal-03677966>

Submitted on 25 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fairness in Selection Problems with Strategic Candidates

VITALII EMELIANOV, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

NICOLAS GAST, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

PATRICK LOISEAU, Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, France

To better understand discriminations and the effect of affirmative actions in selection problems (e.g., college admission or hiring), a recent line of research proposed a model based on *differential variance*. This model assumes that the decision-maker has a noisy estimate of each candidate's quality and puts forward the difference in the noise variances between different demographic groups as a key factor to explain discrimination. The literature on differential variance, however, does not consider the strategic behavior of candidates who can react to the selection procedure to improve their outcome, which is well-known to happen in many domains.

In this paper, we study how the strategic aspect affects fairness in selection problems. We propose to model selection problems with strategic candidates as a contest game: A population of rational candidates compete by *choosing* an effort level to increase their quality. They incur a cost-of-effort but get a (random) quality whose expectation equals the chosen effort. A Bayesian decision-maker observes a noisy estimate of the quality of each candidate (with differential variance) and selects the fraction α of best candidates based on their posterior expected quality; each selected candidate receives a reward S . We characterize the (unique) equilibrium of this game in the different parameters' regimes, both when the decision-maker is unconstrained and when they are constrained to respect the fairness notion of demographic parity. Our results reveal important impacts of the strategic behavior on the discrimination observed at equilibrium and allow us to understand the effect of imposing demographic parity in this context. In particular, we find that, in many cases, the results contrast with the non-strategic setting. We also find that, when the cost-of-effort depends on the demographic group (which is reasonable in many cases), then it entirely governs the observed discrimination (i.e., the noise becomes a second-order effect that does not have any impact on discrimination). Finally we find that imposing demographic parity can sometimes increase the quality of the selection at equilibrium; which surprisingly contrasts with the optimality of the Bayesian decision-maker in the non-strategic case. Our results give a new perspective on fairness in selection problems, relevant in many domains where strategic behavior is a reality.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design**; • **Applied computing** → **Economics**.

Additional Key Words and Phrases: selection problem, fairness, differential variance, strategic candidates

ACM Reference Format:

Vitalii Emelianov, Nicolas Gast, and Patrick Loiseau. 2022. Fairness in Selection Problems with Strategic Candidates. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*, July 11–15, 2022, Boulder, CO, USA. ACM, New York, NY, USA, 29 pages. <https://doi.org/10.1145/3490486.3538287>

1 INTRODUCTION

Selection problems are problems in which, given multiple candidates, a decision-maker selects a fixed fraction of them with the objective of taking the best ones. Selection problems model a number of decision-making situations with high societal implications such as college admission or

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

EC '22, July 11–15, 2022, Boulder, CO, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9150-4/22/07...\$15.00

<https://doi.org/10.1145/3490486.3538287>

hiring. In such cases, it is important to consider the *fairness* of the selection outcome—that is that there is no discrimination against demographic groups defined by *sensitive* attributes. Yet, there is abundant evidence, including in recent years, of discrimination in both college admission and hiring, where the applicants are discriminated by salient demographic attributes such as gender [1, 19], race [5, 17, 30], or age [10].

Recent literature on fairness in selection problem analyzed the problem using models based on two key ingredients to explain discrimination. Kleinberg and Raghavan [21] model selection problems with *implicit bias* (see also [7, 8, 25]), that is, where the decision-maker implicitly under-evaluates the quality of the candidates from disadvantaged demographic groups. On the other hand, Emelianov et al. [12, 13] and Garg et al. [16], following ideas from the economics literature on *statistical discrimination* (see Section 2), assume that the decision maker’s estimate of the candidates quality is unbiased but has a higher variance for some demographic groups (a phenomenon terms *implicit* or *differential variance*). With both types of models, the authors study the discrimination that comes out of baseline decision-makers, and the impact of imposing fairness mechanisms such as the Rooney rule or the four-fifths rule.

In the above literature, the characteristics of the candidates used for selection (in particular, their qualities) are assumed to be fixed and exogenous—i.e., they do not depend on the selection procedure. In practice, however, candidates (i.e., individuals) involved in selection problems can *adapt* to the selection procedure in order to increase the chances of a positive outcome. Such a *strategic behavior* is observed in many domains [28, 34]. A recent thread of literature on *strategic classification* is devoted to modeling and analyzing the impact of this strategic behavior on classification problem [6, 11, 18, 20, 26, 27, 33, 35]. The selection problem, however, is fundamentally different from a classification problem in that the number of positive predictions is constrained—this will in particular lead to competition between individuals, see below. Moreover, this thread of literature did not investigate discrimination issues. This leaves open the key question, which is our focus in this work: *How does the strategic behavior of candidates affect discrimination and the impact of fairness mechanisms in selection problems?*

The selection problem with strategic candidates modeled as a contest game. We propose to model the selection problem with strategic candidates as a contest game (that is, roughly, as a game where candidates compete for a reward). We consider a population of candidates. Each candidate chooses an effort level m that they exert to improve their quality, at a quadratic cost $Cm^2/2$ (where C is a constant coefficient). Each candidate has a latent quality W drawn randomly whose expected value is equal to their selected effort m . A Bayesian decision-maker observes a noisy estimate \hat{W} of the quality W and selects a fraction α of the best candidates based on the posterior expected quality $\tilde{W} = E(W|\hat{W})$. All selected candidates receive a reward of size S , which is a quantitative measure of the benefit that the selection brings to a candidate (e.g., a job position or an education).

In our model, we assume that the population of candidates is divided in two groups: the high-noise (H) and the low-noise (L) group (which refers to the noise in the estimate \hat{W} of the candidates quality). This group-dependent noise represents a form of information inequality common across different demographic groups: decision-makers are less familiar with candidates from certain groups (the H -candidates), such that they are less able to precisely estimate the qualities of these candidates. This phenomenon is called *differential variance* [13, 16]. It is at the basis of the economic theory of statistical discrimination and it was first described by Phelps [29] to explain the racial inequality of wages. In addition to group-dependent noise, we also consider a group-dependent cost coefficient (i.e., $C_H \neq C_L$). The group-dependent cost coefficient models socioeconomic inequalities;

e.g., for students from low-income families it is typically harder to reach a desired level of quality due to costly preparatory courses.

Overview of our results. Our model of a selection problem with strategic candidates defines a (population) game. We first show that this game has a unique Nash equilibrium. Then we focus primarily on characterizing the equilibrium in the regime of large rewards S , which corresponds roughly to high-stake selection problems. In this regime, **we show that at equilibrium, the following discrimination resulting from our model:**

- If the cost coefficient is group-independent ($C_H = C_L$), then the high-noise candidates make a greater effort in average than the low-noise candidates. The latter are underrepresented¹ in the selection, which is counterintuitive and is in contrast with the results in the non-strategic setting studied in [13, 16].
- If the cost coefficient is group-dependent ($C_H \neq C_L$), then the noise level does not affect the outcome of the game. The cost-advantaged candidates (those for whom the cost coefficient is smaller) make a greater average effort compared to the cost-disadvantaged candidates, and the latter are underrepresented, irrespective of the noise. The noise is a second order effect compared to the cost difference, which totally dominates. This offers a potential explanation as to why low-noise candidates are often not underrepresented in practice: this is because the cost of effort for the high-noise candidates is usually large.

In both cases stated above, one of the groups of candidates is always underrepresented. A potential remedy for this is the so-called *demographic parity* mechanism, which imposes that the decision-maker selects candidates of all demographic groups at equal rates. Next, **we characterize the equilibrium when the Bayesian decision-maker is constrained by the demographic parity mechanism** (still in the regime of large rewards S):

- We show that the demographic parity mechanism tends to equalize the average effort of the two groups compared to the unconstrained decision-maker: in most cases, the previously underrepresented group makes a greater average effort and the previously overrepresented group makes a lower average effort.
- We characterize the change in the selection quality (or utility from the decision-maker perspective) from imposing demographic parity. Interestingly, we find that in some cases, the selection quality can improve compared to the unconstrained selection. This is surprising since, in the non-strategic setting, the unconstrained Bayesian decision-maker is optimal [13]. Our result shows that it is no longer true in the strategic setting as in certain cases imposing a fairness mechanism can improve the selection quality, even against a Bayesian baseline. In other cases, we bound the degradation of quality that can result from imposing demographic parity.

For the case of small rewards S , **we get further analytical results on the equilibrium characterization.** We find that the results are different from that of the case of large S , and are similar to the ones obtained in the non-strategic setting [13, 16]: the high-noise candidates are always underrepresented if the selection size α is small enough. This indicates that, if the reward is small, then the impact of the strategic aspect (on discrimination results) is not important. We perform numerical experiments to illustrate the case of intermediate rewards S and how it matches the case of small and large S in their respective regimes.

¹“Underrepresented” in our paper means “having less representation in the selection than its share in the candidates’ population”. This is a classical definition in the algorithmic fairness literature where the notion of demographic parity would mean that the two groups are equally represented.

Finally, we study the convergence of two dynamics to the Nash equilibrium (the best response and fictitious play dynamics). We observe that the trajectories of the best response dynamics converge to limit cycles that have a higher average effort (for both groups) than at equilibrium. The fictitious play dynamics seems to converge to equilibrium in our empirical results.

Implications of our results. Our results show that it is crucial to take into account the strategic nature of candidates involved in selection problem in order to understand discrimination and to predict the effect of imposing a fairness mechanism. They also show that discrimination in selection problems is a somewhat nuanced issue: it depends not only on the strategic aspect of the candidates but also on the range of assumed rewards and on the cost of effort of the different demographic groups. This means that a policy-maker, when considering whether to impose a fairness mechanism in a particular application, should evaluate the population of candidates (their costs of effort, etc.). It is worth noting, also, that we presented our results for a Bayesian decision-maker who computes a posterior estimate of the candidates quality (with knowledge of the group-dependent distributions). Our results are technically easy to extend to a group-oblivious decision-maker sorting the candidates by quality estimate \hat{W} irrespective of their group (another standard baseline); but in that case the high-noise and low-noise groups are reversed (we discuss that in Section 7). Hence, a policy-maker would also need to evaluate the baseline decision-maker they face.

Outline. The rest of the paper is organized as follows. In Section 2, we present the related literature. In Section 3, we formulate the problem in a game-theoretic framework. In Section 4, we show that, for the case of large rewards S , the selection with strategic candidates leads to underrepresentation of one of the two groups of candidates. In Section 5, we study how demographic parity affects the incentives of candidates and the expected quality of the selection for large S . In Section 6, we complement our theoretical results by studying the convergence to equilibria and the case of small and intermediate rewards S . We conclude with a discussion in Section 7.

2 RELATED WORK

Statistical discrimination. The theory of statistical discrimination, initiated by Phelps [29] and Arrow [3], considers the uncertainty of information about individual characteristics to explain racial/gender inequality in decision-making. Phelps [29] develops a model where each individual possesses a latent quality drawn from a fixed group-independent distribution. A Bayesian decision-maker observes a noisy estimate of individual's quality, where the noise is symmetric and zero-mean but has a group-dependent variance. The decision-maker assigns a wage equal to the expected posterior quality. This model is used to explain racial inequalities in wages. Lundberg and Startz [24] extend Phelps' model to a strategic setting by assuming two groups of workers that choose the values of effort according to a group-independent quadratic costs. The effort induces a quality that is assigned randomly but equal to the effort in expectation. The authors show that, in the equilibrium, the high-noise candidates make lower effort and are payed less on average compared to the low-noise candidates; but that if the decision-maker is restricted to not use the group information for wage assignment, then the effort is equal for both groups. In our work, we use a similar model but in the context of selection problems where a fraction of candidates receives a (fixed) reward rather than assigning a variable wage to all candidates as in [24, 29]. We extend the model of Lundberg and Startz [24] to have group-dependent cost-of-effort, which as we shall see has a crucial effect on discrimination and we consider a different fairness mechanism (demographic parity).

Many statistical discrimination models (see e.g., [2, 3, 9] and a survey in [14]) assume that individuals of different groups have identical *a priori* characteristics (e.g., cost-of-effort), but that the decision-maker uses their group-dependent belief when there is imperfect information to assess the performance of individuals in a group. In some cases, the discriminating beliefs (stereotypes) of

decision-makers lead to equilibria in which these stereotypes are fulfilled. In contrast, we consider group-dependent cost-of-effort and group-dependent noise variance. We prove that the game in our model attains a unique equilibrium, and that discrimination occurs due to the group-dependent characteristics—i.e., if they become group-independent in our model, discrimination disappears.

Selection problems in the non-strategic settings. Recent literature investigate statistical discrimination in selection problems, under the name of *differential variance*. Emelianov et al. [13] show that differential variance, in the case of Bayesian decision-maker, leads to the underrepresentation of the high-noise candidates (for selection fractions below 0.5). They also study how quota-based fairness mechanisms (the 80%-rule and the demographic parity) affect the fairness-utility tradeoff. Garg et al. [16] study a similar setting but where the performance of candidates is measured by *multiple* independent and unbiased estimates (from tests). The authors study how affirmative actions and access to testing affect the disparity level and the quality of the selected cohort. Some works also study the selection problem under implicit bias rather than differential variance—i.e., the decision-maker has a non-noisy but biased estimate of the candidates' qualities. Kleinberg and Raghavan [21] show that this type of bias naturally leads to underrepresentation of the disadvantaged group, and they study how a fairness mechanism called the Rooney rule affects the selection quality. This work is extended in [7, 8, 25]. Our work complements those studies by assuming *strategic candidates* who can respond to a policy chosen by a Bayesian decision-maker. Similarly to [16] and [13], we assume that the quality estimates are affected by differential variance, but we do not model bias [21] as we assume a Bayesian decision-maker who can correct for the bias. The main difference is that we assume that the quality distribution is not fixed but chosen by candidates to maximize their payoff which equals to the expected reward minus the cost-of-effort.

Fairness in contests. A classical model of contest is given by Lazear and Rosen [22]: individuals make a costly effort m that induces a noisy quality $W = f(m) + \varepsilon$, where f is an increasing function of the effort; the player having the largest quality wins the prize S . Fairness in contest games was studied from different perspectives in economics and computer science literature [15]. Schotter and Weigelt [32] assume a two-player contest with a cost-advantaged and a cost-disadvantaged player—each pays a quadratic cost-of-effort but with different coefficients. They show that, at equilibrium, the cost-disadvantaged player makes a lower effort. Then they show that an affirmative action (adding a bonus to the score of the cost-disadvantaged player) leads all players to make lower efforts but increases the winning probability of the cost-disadvantaged player. Our model shares similarities with that of [32] (in particular group-dependent quadratic costs), but also has important differences: we consider an infinite population of candidates and we include group-dependent noise. Our results are also different as we analyze the dependence on S and we study another fairness mechanism (demographic parity)—in particular we find that it increases the effort of previously underrepresented group and sometimes even of both groups.

Strategic ranking. Following the strategic classification literature (see previous section), Liu et al. [23] study the *strategic ranking* problem. They assume that each individual has a latent ranking and makes some costly effort to affect it. The effort induces a score $g(\text{effort}) \cdot f(\text{latent fixed rank})$, for some fixed strictly increasing functions f and g , which in turn results in a post-effort ranking. Liu et al. [23] study how different ranking reward functions (with and without randomization) of a fixed selection capacity c affect the characteristics of individuals and of the selection at equilibrium (average welfare and scores). Finally, they consider a setting with two groups of candidates that differ in a multiplicative parameter $\gamma > 1$ affecting the score and study the welfare gap for groups as a function of c . In our model, we do not assume any pre-effort ranking—the score of an individual is purely determined by the effort and by a group-dependent noise. Our model is simpler compared

to [23] and is designed to capture the effects of the group-dependent cost-of-effort and noise in selection problems. This allows us to state concrete results involving the parameters of interest (the cost coefficient and noise variance). We also consider how a fairness mechanism (demographic parity) affects the group representations and the quality of selection at equilibrium.

3 THE MODEL

3.1 Candidates and decision-maker

Candidates model. We assume a non-atomic game with a unit mass of candidates indexed by $i \in [0, 1]$. There are two groups of candidates: H and L . The letter $G \in \{H, L\}$ will denote any of these two groups, and the proportion of candidates from group G is $p_G \in (0, 1)$. Each candidate i has a quality that depends on the *effort* m_i that this candidate makes. In college admission or in hiring, the effort m_i of a candidate i can be interpreted as candidate's achievements. It might, for instance, represent the number of courses followed by a student, the quality of number of degrees obtained, etc. We assume that a candidate i that chooses to make an effort $m_i \geq 0$ has a quality W_i that is normally distributed with mean m_i and variance η^2 :

$$W_i \sim \mathcal{N}(m_i, \eta^2).$$

Making an effort m costs a candidate from group G a quadratic cost $C_G m^2/2$. The population-dependent cost coefficient C_G can model socioeconomic factors like the income of the parents or the country of origin; these factors might make harder for some candidates than others to make the same effort m .

Decision-maker. A decision-maker wants to select the candidates having the largest qualities. To do so, the decision-maker observes a noisy estimate \hat{W}_i of the quality W_i of each candidate i :

$$\hat{W}_i = W_i + \hat{\sigma}_{G_i} \cdot \varepsilon_i,$$

where ε_i is a zero-mean centered random variable from the standard normal distribution $\mathcal{N}(0, 1)$; the noise variance $\hat{\sigma}_{G_i}^2$ is assumed group-dependent.² The quality estimate \hat{W}_i is a noisy but unbiased measurement of the quality W_i of a candidate i , e.g., an interview result. The group-dependent variance of noise $\hat{\sigma}_G^2$ models the information inequality: if interviewers are more familiar with candidates of some demographic groups, they are more confident in their evaluation compared to that of candidates of other groups. Without loss of generality, we assume that $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$. We, thus, refer to H -candidates as *high-noise* candidates, and we refer to L -candidates as *low-noise* candidates.

We assume that the decision-maker knows³ the effort m of all candidates, as well as the variances η^2 and $\hat{\sigma}_G^2$, and selects a proportion $\alpha \in (0, 1)$ of the candidates. The decision-maker aims at maximizing the expected quality of selected candidates and, therefore, selects the α proportion having the largest expected quality \tilde{W} . Using the property of conditional expectation for bivariate normal random variables, we can write the expectation of quality W_i given \hat{W}_i as

$$\tilde{W}_i = E(W_i | \hat{W}_i) = \hat{W}_i \rho_{G_i}^2 + (1 - \rho_{G_i}^2) m_i, \quad (1)$$

where $\rho_{G_i} = \eta / \sqrt{\eta^2 + \hat{\sigma}_{G_i}^2} \in [0, 1]$ is the correlation coefficient between the quality W_i and its estimate \hat{W}_i . Since \tilde{W}_i is a linear function of \hat{W}_i , and \hat{W}_i is distributed normally, the expected quality \tilde{W}_i follows a normal distribution with mean m_i and variance $\tilde{\sigma}_{G_i}^2 = \eta^4 / (\hat{\sigma}_{G_i}^2 + \eta^2)$. Note that the larger the value of noise $\hat{\sigma}_{G_i}^2$, the more the values of \tilde{W}_i are concentrated around its mean value

²Note that, for simplicity, we assume that the variance of the quality η^2 does not depend on the candidate's group; Our results can be extended to the case of group-dependent η_G^2 (see Section 7).

³Our results can be extended to the case where the effort m is not observable (see Section 7).

m_i (the smaller the variance $\tilde{\sigma}_{G_i}^2$). From (1), we observe that the decision-maker puts a higher weight on the effort m for the high-noise candidates compared to the low-noise candidates: for the same level of effort, the high-noise candidates will be seen by the decision-maker as having less variability of expected quality compared to that of the low-noise candidates.

3.2 The population game

As the decision-maker selects the candidates having the largest expected quality, it will select all candidates whose expected quality \tilde{W} is larger than some selection threshold θ . We denote by $x_G(m; \theta)$ the probability for a candidate of group G to be selected when their effort is m and the selection threshold is θ . It equals:⁴

$$x_G(m; \theta) = \mathbb{P}(\tilde{W} \geq \theta) = \Phi\left(\frac{m - \theta}{\tilde{\sigma}_G}\right), \quad (2)$$

where Φ is the cumulative distribution function of the standard normal distribution $\mathcal{N}(0, 1)$.

We assume that each selected candidate receives a positive reward S , whereas the candidates who are not selected get a reward of 0. Hence, the *payoff* u_G of a candidate from population G is

$$u_G(m; \theta) = S \cdot x_G(m; \theta) - C_G \cdot m^2/2. \quad (3)$$

Given a selection threshold θ , each candidate strategically decides on the effort m that maximizes their expected payoff $u_G(m; \theta)$. Following the classical definition [31], we call it a *pure best response* and we denote the set of all *pure best response* strategies of a candidate as

$$b_G(\theta) = \{m \text{ such that } \forall m' \geq 0 : u_G(m; \theta) \geq u_G(m'; \theta)\}.$$

We say that the best response is unique if the set of best responses $b_G(\theta)$ is reduced to a singleton. In such a case, by abuse of notation, we denote by $b_G(\theta)$ the unique best response.

Similarly, by $\beta_G(\theta)$ we denote the set of all *mixed best responses* of a candidate. It is the set of all probability distributions over the set of pure best responses $b_G(\theta)$.

Selection threshold. For each population $G \in \{H, L\}$, we denote by μ_G the distribution of efforts used by this population G . It is a probability distribution⁵ on \mathbf{R}^+ . We denote by $\boldsymbol{\mu} = (\mu_H, \mu_L)$ the effort distributions of the two populations. We denote the cumulative distribution function of the expected quality \tilde{W} induced by μ_G by F_{μ_G} , and the cumulative distribution function of expected quality of the total population by $F_{\boldsymbol{\mu}} = p_H F_{\mu_H} + p_L F_{\mu_L}$. The decision-maker selects the α -best candidates. Hence, it will select all candidate whose expected quality \tilde{W} is above the $(1 - \alpha)$ -quantile of the distribution $F_{\boldsymbol{\mu}}$, that we denote by $\theta(\boldsymbol{\mu}) = F_{\boldsymbol{\mu}}^{-1}(1 - \alpha)$.

The game. The above definitions describe a population game between the candidates. We denote the game by \mathcal{G}^{un} , where the superscript “un” emphasizes that the decision-maker is *unconstrained*,⁶ i.e., it selects the best α candidates based on the expected quality \tilde{W} . Note that the game \mathcal{G}^{un} is parameterized by the reward size S , the variance of quality η^2 , the noise variances $\hat{\sigma}_G^2$, the cost coefficients C_G , and the selection size α .

3.3 Existence and uniqueness of the Nash equilibrium

We use the standard definition of Nash equilibrium for populations games [31]:

⁴In the rest of the paper, to simplify the notation, we will drop the index i .

⁵If all candidates of the population G make the same effort, μ_G is a *pure strategy*. Otherwise, μ_G is a *mixed strategy*.

⁶In Section 5, we study a decision-maker faced with a fairness constraint (called demographic parity).

Table 1. Summary of notation.

$W_i \sim \mathcal{N}(m_i, \eta^2)$	quality of candidate i
$\hat{W}_i = W_i + \varepsilon_i \hat{\sigma}_{G_i}$	quality estimate of candidate i ($\varepsilon_i \sim \mathcal{N}(0, 1)$)
$\tilde{W}_i = \mathbb{E}(W_i \hat{W}_i)$	expected quality of candidate i
μ_G	distribution of effort for the population G
$\boldsymbol{\mu} = (\mu_H, \mu_L)$	distribution of effort for the total population
F_{μ_G}	CDF of the expected quality \tilde{W} for the population G
$F_{\boldsymbol{\mu}} = p_H F_{\mu_H} + p_L F_{\mu_L}$	CDF of the expected quality \tilde{W} for the total population
$\alpha \in (0, 1)$	selection size
$\theta(\boldsymbol{\mu}) = F_{\boldsymbol{\mu}}^{-1}(1 - \alpha)$	selection threshold
p_G	proportion of G -candidates in the total population
$x_G(m; \theta)$	selection probability for a G -candidate with effort m given the threshold θ

DEFINITION 1 (NASH EQUILIBRIUM). A pair of effort distributions $\boldsymbol{\mu} = (\mu_H, \mu_L)$ is called an equilibrium of the game \mathcal{G}^{un} if for all populations $G \in \{H, L\}$, the support of μ_G is included in the set of best responses $b_G(\theta(\boldsymbol{\mu}))$, where $\theta(\boldsymbol{\mu}) = F_{\boldsymbol{\mu}}^{-1}(1 - \alpha)$.

A Nash equilibrium is a situation where no candidate has an incentive to change its decision: if the population plays $\boldsymbol{\mu}$, then the selection threshold will be $\theta(\boldsymbol{\mu})$. Hence, a candidate of a group G does not have an incentive to play a strategy that is not in $b(\theta(\boldsymbol{\mu}))$. As the support of μ_G is included in $b(\theta(\boldsymbol{\mu}))$, this implies that all candidates cannot obtain a higher payoff by unilaterally changing their decision.

In the rest of the paper, we will study the property of the Nash equilibrium of the game, which exists and is unique as guaranteed by the following theorem.

THEOREM 1. The game \mathcal{G}^{un} has a unique Nash equilibrium, that we denote by $\boldsymbol{\mu}^{\text{un}} = (\mu_H^{\text{un}}, \mu_L^{\text{un}})$.

We will also denote by $\theta^{\text{un}} = F_{\boldsymbol{\mu}^{\text{un}}}^{-1}(1 - \alpha)$ the selection threshold at equilibrium. Note that the equilibrium and the selection threshold depend on the game parameters, including the reward S . In the following, we will study the properties of the equilibrium as S grows. For simplicity of exposition, we omit the dependency on S and write $\boldsymbol{\mu}^{\text{un}}$ or θ^{un} instead of $\boldsymbol{\mu}^{\text{un}}(S)$ or $\theta^{\text{un}}(S)$.

We provide a detailed proof of Theorem 1 in Appendix A.4 whose main ingredient is to show that there is a one-to-one mapping between the equilibria of the game and the fixed points of a multi-valued function T :

$$T(\theta) = \{F_{\boldsymbol{\mu}}^{-1}(1 - \alpha) : \mu_G \in \beta_G(\theta)\}. \quad (4)$$

This simplifies the problem, as now we need to study the fixed points of a function of a single variable. By studying the first derivative of T , we show that the function T has a unique fixed point and, hence, the defined game has also a unique equilibrium.

3.4 Summary of main notation

We summarize the main notation in Table 1. Recall that, without loss of generality, we assume that $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$ (which also implies that $\tilde{\sigma}_H^2 < \tilde{\sigma}_L^2$). We, thus, refer to H -candidates as *high-noise*, and we refer to L -candidates as *low-noise*. To refer to the group with higher/lower cost-of-effort, we use the names *cost-disadvantaged/cost-advantaged*. If, for example, $C_H > C_L$, then we say that H -candidates are *cost-disadvantaged* and that L -candidates are *cost-advantaged*.

4 EQUILIBRIUM CHARACTERIZATION AND RESULTING DISCRIMINATION

In this section, we characterize the equilibrium of the game for large⁷ rewards S . We consider the case of large S as it models the competition in selection procedures with high stakes, e.g., hiring CEOs or college admission to high-ranked schools.

4.1 Properties of the best response

We start with the characterization of the best response $b_G(\theta)$ of a candidate. In Lemma 1 (whose proof is deferred to Appendix A.2), we show that, when S is large, the best response $b_G(\theta)$ is unique for all thresholds θ , except one, that we call a *dropout threshold* θ_G^d . We show that the value of the best response $b_G(\theta)$ increases with the threshold θ until the latter reaches θ_G^d ; it then drops down and decreases when $\theta \geq \theta_G^d$. This means that when the selection threshold θ is too high, the candidates lose incentives to make any effort. To emphasize the dependency of the dropout threshold on the reward S , we write $\theta_G^d(S)$.

LEMMA 1 (BEST RESPONSE FOR LARGE S). *There exists S_0 such that for all rewards $S \geq S_0$, there exists a threshold $\theta_G^d(S)$, called the dropout threshold, such:*

- (i) *When the selection threshold is $\theta = \theta_G^d(S)$, there are two pure best response strategies: $b_G(\theta_G^d(S)) = \{b_G^{\min}(\theta_G^d(S)), b_G^{\max}(\theta_G^d(S))\}$. They satisfy:*

$$\lim_{S \rightarrow \infty} b_G^{\min}(\theta_G^d(S)) = 0, \quad \lim_{S \rightarrow \infty} \frac{b_G^{\max}(\theta_G^d(S))}{\theta_G^d(S)} = 1.$$

- (ii) *For all $\theta \neq \theta_G^d(S)$, the pure best response $b_G(\theta)$ is unique. Moreover, for any $\gamma \in (0, 1)$, we have*

$$\lim_{S \rightarrow \infty} b_G(\theta_G^d(S)/\gamma) = 0 \text{ and } \lim_{S \rightarrow \infty} \frac{b_G(\gamma\theta_G^d(S))}{\theta_G^d(S)} = \gamma.$$

Using the notation for asymptotic equivalence, we can write the statement of the theorem in a simpler form. For example, from the second part of (ii), we can infer that $b_G(\gamma\theta_G^d) \sim \gamma\theta_G^d$ for any $\gamma \in (0, 1)$. Therefore, the above lemma shows that for a given candidate, the best response $b_G(\theta)$ increases essentially linearly up to the dropout threshold θ_G^d , and then drops to 0 when the selection threshold is too high, i.e., $\theta > \theta_G^d$. We will illustrate this later in Fig. 1.

Note that the dropout threshold θ_G^d is group-dependent. As we will see later in Section 4.3, the most represented group at equilibrium will be the one having the largest dropout threshold. Hence, in Lemma 2, we show the asymptotic behavior of the dropout threshold as a function of S . As we expect, the dropout threshold increases with the reward S , and decreases with the cost coefficient C_G . A less intuitive property is its relation with the noise variance $\hat{\sigma}_G^2$: we show that the dropout threshold increases as the noise variance $\hat{\sigma}_G^2$ increases. Lemma 2 is proven in Appendix A.3.

LEMMA 2 (DROPOUT THRESHOLD FOR LARGE S). *Let $\theta_G^d(S)$ be the dropout threshold.*

- (i) *For any C_G and G , the dropout threshold can be expressed as*

$$\theta_G^d(S) = \sqrt{2S/C_G} (1 + o(1)) \text{ as } S \rightarrow \infty.$$

- (ii) *If $C_H = C_L$ and $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$, then there exists S_0 such that for all $S \geq S_0$, $\theta_H^d(S) > \theta_L^d(S)$.*

From (i), we conclude that the dependency of the dropout threshold θ^d on the cost coefficients C_G is of higher order than the dependency on the noise variance $\hat{\sigma}_G^2$. In other words, when $C_H > C_L$, we have $\theta_H^d(S) < \theta_L^d(S)$ for large enough S , regardless of the noise variance $\hat{\sigma}_G^2$. When both groups have

⁷In Section 6.2, we show theoretical and numerical results for small and intermediate rewards S .

equal cost coefficients ($C_H = C_L$), the noise variance matters, and $\theta_H^d(S) > \theta_L^d(S)$ when $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$ (or, equivalently, when $\tilde{\sigma}_H^2 < \tilde{\sigma}_L^2$).

4.2 Equilibrium strategy

As we show in the previous section, the dropout threshold θ_G^d is an important characteristics of candidates' best response: for thresholds θ smaller than the dropout threshold θ_G^d , the candidates make an effort proportional to the threshold whereas for thresholds larger than the dropout, the candidates make nearly zero efforts. In Theorem 2 given below, we prove that, for large rewards, the selection threshold at equilibrium is equal to the dropout threshold of one of the two groups. Note that this theorem is not stated in terms of the groups H and L but in terms of groups G_1 and G_2 , where G_1 is the group that has the largest dropout threshold, which holds if $C_1 < C_2$ or ($C_1 = C_2$ and $\hat{\sigma}_1^2 > \hat{\sigma}_2^2$). Hence, the population G_1 can correspond to the population L if $C_H > C_L$ (it corresponds, otherwise, to the population H as we assumed that $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$).

THEOREM 2 (EQUILIBRIUM STRATEGIES). *Fix $\alpha \in (0, 1)$ and two populations of candidates, G_1 and G_2 , such that $(C_1 < C_2)$ or $(C_1 = C_2$ and $\hat{\sigma}_1^2 > \hat{\sigma}_2^2)$. Then, there exists a reward S_0 such that for $S \geq S_0$:*

- (i) *If $\alpha \leq p_1$, then the equilibrium threshold $\theta^{\text{un}}(S)$ is $\theta_1^d(S)$. In this case:*
 - *The G_1 -candidates play a mixed strategy that consists in playing b_1^{max} with probability τ_1 and b_1^{min} , with probability $1 - \tau_1$, where $\lim_{S \rightarrow \infty} \tau_1 = \alpha/p_1$.*
 - *The G_2 -candidates play the pure strategy $b_2(\theta_1^d(S))$ and $\lim_{S \rightarrow \infty} b_2(\theta_1^d(S)) = 0$.*
- (ii) *If $\alpha > p_1$, then the equilibrium threshold $\theta^{\text{un}}(S)$ is $\theta_2^d(S)$. In this case:*
 - *The G_1 -candidates play the pure strategy $b_1(\theta_2^d(S))$ and $\lim_{S \rightarrow \infty} b_1(\theta_2^d(S))/\theta_2^d(S) = 1$.*
 - *The G_2 -candidates play a mixed strategy that consists in playing b_2^{max} with probability τ_2 and b_2^{min} with probability $1 - \tau_2$, where $\lim_{S \rightarrow \infty} \tau_2 = (\alpha - p_1)/p_2$.*

PROOF SKETCH. To simplify notation, we omit the dependence on S for all variables. Let us prove that the dropout threshold θ_1^d is the fixed point of T for the case (i). We specify the efforts made in response to θ_1^d , and we show that they lead to the same selection threshold θ_1^d , i.e., θ_1^d is the fixed point of T . The case (ii) can be proven similarly; we provide the complete proof in Appendix A.5.

We can show that, as $S \rightarrow \infty$, we have $x_G^{\text{max}} := x_G(b_G^{\text{max}}(\theta_G^d); \theta_G^d) \xrightarrow{S \rightarrow \infty} 1$ and $x_G^{\text{min}} := x_G(b_G^{\text{min}}(\theta_G^d); \theta_G^d) \xrightarrow{S \rightarrow \infty} 0$ for all $G \in \{G_1, G_2\}$.

If $\alpha \leq p_1$, then, assume that G_1 -candidates randomize their strategy by playing b_1^{max} with the probability τ_1 and b_1^{min} with the probability $1 - \tau_1$. The G_2 -candidates play the deterministic strategy $b_2(\theta_1^d) \xrightarrow{S \rightarrow \infty} 0$. The probability τ_1 can be found from the budget constraint: $\alpha = p_1(\tau_1 x_1^{\text{max}} + (1 - \tau_1)x_1^{\text{min}}) + p_2 x_2(\theta_1^d, b_2(\theta_1^d)) \xrightarrow{S \rightarrow \infty} p_1 \tau_1$, so $\tau_1 \xrightarrow{S \rightarrow \infty} \alpha/p_1$. The defined strategies satisfy the budget constraint, so θ_1^d is the fixed point of T and, hence, the defined distribution of effort is the equilibrium of the game \mathcal{G}^{un} . \square

Fig. 1 illustrates the results of Theorem 2. In Fig. 1a, we show the case of group-dependent noise but group-independent cost coefficient ($C_H = C_L$). In this case, the H -candidates have a higher dropout threshold compared to the L -candidates, hence, for $\theta_L^d < \theta < \theta_H^d$, the H -candidates make a non-zero effort while the L -candidates make a nearly-zero effort. In our illustration, the selection size $\alpha = 0.1$ is smaller than the proportion of H -candidates $p_H = 0.5$. We can verify that for $\theta = \theta_H^d$, if a proportion α/p_H of H -candidates plays b_H^{max} , and the rest of H -candidates plays b_H^{min} , then such a strategy satisfies the budget constraint. Hence, θ_H^d is the fixed point of the function T , so $\theta^{\text{un}} = \theta_H^d$.

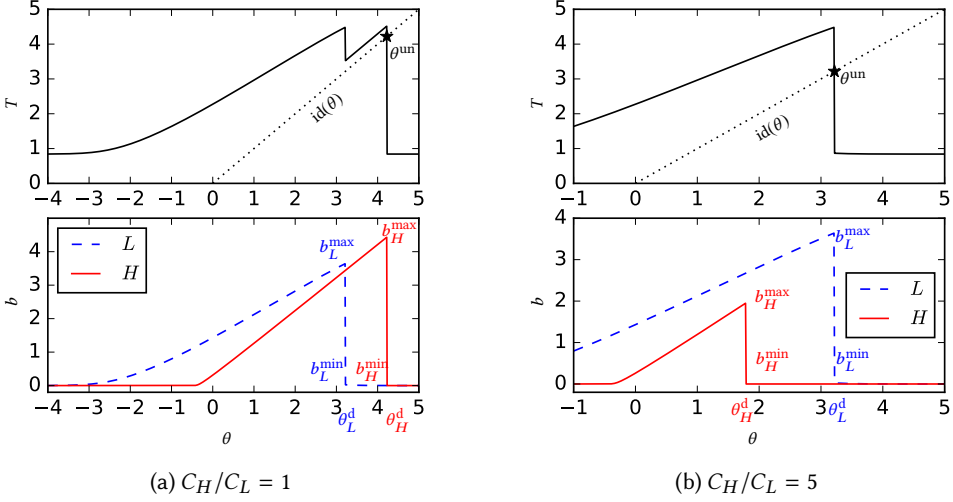


Fig. 1. Best response functions and the Nash equilibrium for $S = 10$, $C_L = 1$, $\tilde{\sigma}_H = 0.1$ and $\tilde{\sigma}_L = 1$, $p_H = 0.5$ and $\alpha = 0.1$. Both figures illustrate the case (i) of Theorem 2. The dotted line on the uppermost panels is the identity function $\text{id}(\theta) = \theta$.

In Fig. 1b, we illustrate the case of group-dependent cost coefficient ($C_H > C_L$). In this case, the L -candidates have a higher dropout threshold compared to the H -candidates, hence, for $\theta_H^d < \theta \leq \theta_L^d$, the L -candidates make a non-zero effort while H -candidates make a nearly-zero effort. For the purpose of illustration, we again assume that the selection size $\alpha = 0.1$ is smaller than the proportion of L -candidates ($p_L = 0.5$). Similarly to the previous case, we can verify that for $\theta = \theta_L^d$, if a proportion α/p_L of L -candidates plays b_L^{\max} , and the rest of L -candidates plays b_L^{\min} , then this strategy satisfies the budget constraint, so $\theta^{\text{un}} = \theta_L^d$.

4.3 Discrimination due to the group-dependent variance and cost

In Theorem 2, we show that the equilibrium distribution of effort μ^{un} depends on the relation between the dropout thresholds θ_G^d for different populations, H and L . We now study how the noise variance $\hat{\sigma}_G^2$ and the cost coefficient C_G affect the representation of groups in the selection at equilibrium. Overall, we show that the group-dependent noise variance and the group-dependent cost coefficient can both lead to underrepresentation of some group of candidates at equilibrium.

For a given population G , we denote by $\bar{\mu}_G^{\text{un}}$ the *average effort* that candidates from group G exert at equilibrium, and we denote by \bar{x}_G^{un} the *selection rate* which is the probability for a randomly chosen candidate from group G to be selected. Since $x_G(m; \theta)$ is the probability for a G -candidate to be selected when exerting an effort m (see (2)), these quantities satisfy:

$$\bar{\mu}_G^{\text{un}} = \int_{m \geq 0} m d\mu_G^{\text{un}}(m); \quad \bar{x}_G^{\text{un}} = \int_{m \geq 0} x_G(m; \theta^{\text{un}}) d\mu_G^{\text{un}}(m).$$

We say that the group H (or L) is *underrepresented* if $\bar{x}_H < \bar{x}_L$ (or $\bar{x}_L < \bar{x}_H$). Note that when we say ‘underrepresented’, it means that a demographic group has less representation in the selection than its share in the population of all candidates. This definition is not conditioned on the assumption that the mean of the qualities W is the same for both groups, but this assumption does make the notion of demographic parity more obviously appealing since there is no a priori distinction in the

average quality between the groups. Nevertheless, we do not claim that demographic parity would be justified only under this assumption.

In Theorem 3 below, we show that if the cost coefficient C_G is group-independent ($C_H = C_L$), then the high-noise H -candidates make a larger effort compared to that of the low-noise L -candidates, and the latter are underrepresented. However, if the cost coefficient is group-dependent ($C_H \neq C_L$), the noise variance does not play a role if the reward S is large enough: the cost-advantaged candidates make larger effort compared to that of the cost-disadvantaged candidates, and as a result, the latter are underrepresented. This theorem also shows that, as S goes to infinity, the ratios of efforts and selection rates can grow unbounded. For example, in the case (i) of Theorem 3, if there are enough H -candidates to fill the selection budget α , then L -candidates will asymptotically not be selected when S goes to infinity. We will see in Section 6.2, that for moderate values of reward S , the L -candidates still have some representation in the selection, but that can be very small.

THEOREM 3 (DISCRIMINATION FOR LARGE REWARDS S). *Let μ_H^{un} and μ_L^{un} , x_H^{un} and x_L^{un} be the equilibrium effort distributions and selection rates of the game \mathcal{G}^{un} .*

- (i) *If $C_H < C_L$ or if $C_H = C_L$ and $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$, then there exists S_0 such that for all $S \geq S_0$ the high-noise H -candidates make greater effort compared to the low-noise L -candidates, and the latter are underrepresented:*

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_L^{\text{un}}}{\bar{\mu}_H^{\text{un}}} = \lim_{S \rightarrow \infty} \frac{\bar{x}_L^{\text{un}}}{\bar{x}_H^{\text{un}}} = \begin{cases} 0 & \text{if } p_H \geq \alpha, \\ \frac{\alpha - p_H}{1 - p_H} < 1 & \text{if } p_H < \alpha. \end{cases}$$

- (ii) *If $C_H > C_L$, then there exists S_0 such that for all $S \geq S_0$ the cost-advantaged L -candidates make a greater effort compared to the cost-disadvantaged H -candidates, and the latter are underrepresented:*

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{un}}}{\bar{\mu}_L^{\text{un}}} = \lim_{S \rightarrow \infty} \frac{\bar{x}_H^{\text{un}}}{\bar{x}_L^{\text{un}}} = \begin{cases} 0 & \text{if } p_L \geq \alpha, \\ \frac{\alpha - p_L}{1 - p_L} < 1 & \text{if } p_L < \alpha. \end{cases}$$

PROOF. To prove (i), we compute the average effort and the selection rate given the equilibrium mixed strategies defined in Theorem 2. We consider the case of $\alpha \leq p_H$ and the case of $\alpha > p_H$ separately.

If $\alpha \leq p_H$, then $\bar{\mu}_H^{\text{un}} = b_H^{\text{max}} \tau_H + b_H^{\text{min}} (1 - \tau_H) = \alpha / p_H \cdot \sqrt{2S / C_H} (1 + o(1))$, $S \rightarrow \infty$ and $\bar{\mu}_L^{\text{un}} = b_L(\theta_H^d) \xrightarrow{S \rightarrow \infty} 0$, so $\bar{\mu}_L^{\text{un}} / \bar{\mu}_H^{\text{un}} \xrightarrow{S \rightarrow \infty} 0$. Similarly, the expected selection rates for H - and L -candidates:

$$\begin{aligned} \bar{x}_H^{\text{un}} &= x_H^{\text{max}} \tau_H + x_H^{\text{min}} (1 - \tau_H) \xrightarrow{S \rightarrow \infty} \alpha / p_H, \\ \bar{x}_L^{\text{un}} &= x_L(\theta_H^d, b_L(\theta_H^d)) \xrightarrow{S \rightarrow \infty} 0, \end{aligned}$$

which implies that $\bar{x}_L^{\text{un}} / \bar{x}_H^{\text{un}} \xrightarrow{S \rightarrow \infty} 0$.

If $\alpha > p_H$, then we, again, consider the strategies in the equilibrium:

$$\begin{aligned} \bar{\mu}_H^{\text{un}} &= b_H(\theta_L^d) = \sqrt{2S / C_L} (1 + o(1)), \quad S \rightarrow \infty, \\ \bar{\mu}_L^{\text{un}} &= b_L^{\text{max}} \tau_L + b_L^{\text{min}} (1 - \tau_L) = (\alpha - p_H) / p_L \cdot \sqrt{2S / C_L} (1 + o(1)), \quad S \rightarrow \infty. \end{aligned}$$

Similarly, for the selection rates, we have: $\bar{x}_H^{\text{un}} = x_H(\theta_L^d, b_H(\theta_L^d)) \xrightarrow{S \rightarrow \infty} 1$, and $\bar{x}_L^{\text{un}} = x_L^{\text{max}} \tau_L + x_L^{\text{min}} (1 - \tau_L) \xrightarrow{S \rightarrow \infty} (\alpha - p_H) / p_L$. The proof of (ii) is identical to the proof of (i) so we omit it. \square

5 EFFECTS OF THE DEMOGRAPHIC PARITY MECHANISM ON THE SELECTION

In Theorem 3, we show that the equilibrium in the game \mathcal{G}^{un} leads to underrepresentation of one of the groups of candidates: for the group-independent cost coefficient ($C_H = C_L$), the low-noise candidates are underrepresented; for the group-dependent cost coefficient ($C_H \neq C_L$), the cost-disadvantaged candidates are underrepresented.

To reduce the inequality of representation, colleges (and employers) sometimes perform *affirmative actions*. This can take the form of quotas for low-income or minority groups or any forms of promotions. Among those affirmative actions are ones that make sure that the selection rates for the groups are close, meaning that the proportions of both groups in the selection should be close to proportions of the groups in the total population. The strongest condition among those is the *demographic parity* (see [4]) which requires that selection rates for both populations must be equal: $\bar{x}_H = \bar{x}_L$. This implies, given the budget constraint $\bar{x}_H p_H + \bar{x}_L p_L = \alpha$, that $\bar{x}_H = \bar{x}_L = \alpha$.

5.1 The game with the demographic parity mechanism

The demographic parity mechanism removes the competition among the groups of candidates, H and L , as from each group G a proportion $\alpha \in (0, 1)$ must be accepted. Hence, the game with the demographic parity mechanism, that we denote by \mathcal{G}^{dp} , can be represented as two independent games, $\mathcal{G}_H^{\text{un}}$ and $\mathcal{G}_L^{\text{un}}$, each defined for a single group, H or L . Note that for this game, the selection thresholds will be different for the two groups: for a group G , the decision-maker will select all candidates whose expected quality is greater or equal than $\theta(\mu_G) = F_{\mu_G}^{-1}(1 - \alpha)$, which corresponds to the α best fraction of this group—it is the quantile of the distribution of expected qualities induced by μ_G and not the one induced by μ as in the unconstrained case.

The equilibrium of the game \mathcal{G}^{dp} where the Bayesian decision-maker has a demographic parity constraint is defined as follows.

DEFINITION 2 (NASH EQUILIBRIUM OF THE GAME \mathcal{G}^{dp}). *A pair of effort distributions $\mu = (\mu_H, \mu_L)$ is an equilibrium of the game with the demographic parity constraint \mathcal{G}^{dp} if for both groups $G \in \{H, L\}$:*

$$\mu_G \in \beta_G(\theta(\mu_G)), \text{ where } \theta(\mu_G) = F_{\mu_G}^{-1}(1 - \alpha).$$

Mimicking the unconstrained case, we denote by $\mu^{\text{dp}} = (\mu_H^{\text{dp}}, \mu_L^{\text{dp}})$ the equilibrium of this game, and by $\theta_G^{\text{dp}} = F_{\mu_G}^{-1}(1 - \alpha)$ the group-dependent equilibrium selection threshold. The superscript “dp” emphasizes that the decision-maker is *demographic parity-constrained*.

Since the game \mathcal{G}^{dp} can be represented as two separate games, $\mathcal{G}_H^{\text{un}}$ and $\mathcal{G}_L^{\text{un}}$, that has a unique Nash equilibrium according to Theorem 1, the game \mathcal{G}^{dp} also has a unique Nash equilibrium.

5.2 Efforts induced by the demographic parity mechanism

In the previous section, we showed that, for large rewards S , it is possible that only one of the two groups makes a positive effort while the other group considers that the game is not worth playing because of a too unfair competition. The situation is radically different with demographic parity mechanism as each candidate competes with similar candidates. As we show below, the demographic parity mechanism pushes the previously underrepresented group to make more effort than before. Moreover, it can also push the previously overrepresented group to make more effort.

In the first theorem below, we characterize the equilibrium strategy of the game with the demographic parity mechanism \mathcal{G}^{dp} . This result is a direct corollary of Theorem 2 as we consider two separated games with an unconstrained decision-maker.

THEOREM 4 (EQUILIBRIUM STRATEGY WITH THE DEMOGRAPHIC PARITY MECHANISM). *There exists S_0 such that for all $S \geq S_0$, the equilibrium of the game with the demographic parity mechanism \mathcal{G}^{dp}*

is a pair of distributions $\mu^{\text{dp}} = (\mu_H^{\text{dp}}, \mu_L^{\text{dp}})$ where for each group $G \in \{H, L\}$, μ_G^{dp} consists in playing b_G^{max} with probability τ_G and b_G^{min} with probability $1 - \tau_G$, where $\lim_{S \rightarrow \infty} \tau_G = \alpha$.

PROOF. According to the demographic parity mechanism, the selection rate per each group G must be equal to $\bar{x}_G = \alpha$. Hence, this theorem can be seen as a special case of Theorem 2 but with a single group of mass 1 out of which we need to select the best $\alpha \in (0, 1)$. By applying directly the result of Theorem 2, we show that the proposed strategy is the equilibrium strategy. \square

In Corollary 1 given below, we compare the efforts made by two groups at equilibrium. We show that, if the cost coefficient is group-independent, then the demographic parity mechanism equalizes the effort as S grows (together with the selection rates as $\bar{x}_H^{\text{dp}} = \bar{x}_L^{\text{dp}}$ by definition). For group-dependent cost coefficient, the cost-disadvantaged H -candidates make lower average effort compared to that of L -candidates yet the average effort ratio is bounded by $\sqrt{C_H/C_L}$. This is in contrast to Section 4.3 where we show that the average effort ratio can be unbounded in the case of the unconstrained decision-maker.

COROLLARY 1 (EQUILIBRIUM EFFORT RATIO IN \mathcal{G}^{dp}). *Let μ_H^{dp} and μ_L^{dp} be the equilibrium effort distributions in the game with the demographic parity mechanism \mathcal{G}^{dp} . The average efforts of both populations $G \in \{H, L\}$ satisfy:*

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{dp}}}{\bar{\mu}_L^{\text{dp}}} = \sqrt{\frac{C_L}{C_H}}.$$

In particular, if $C_H = C_L$, then the average efforts of both populations grow at equal rates. If $C_H \neq C_L$, then the cost-disadvantaged candidates make a lower average effort compared to that of the cost-advantaged candidates.

PROOF. Using the equilibrium strategies of the game with the demographic parity mechanism \mathcal{G}^{dp} from Theorem 4, we show that: $\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{dp}}}{\bar{\mu}_L^{\text{dp}}} = \lim_{S \rightarrow \infty} \frac{\tau_H b_H^{\text{max}} + (1 - \tau_H) b_H^{\text{min}}}{\tau_L b_L^{\text{max}} + (1 - \tau_L) b_L^{\text{min}}} = \sqrt{\frac{C_L}{C_H}}$. \square

By reducing the competition between groups, affirmative action policies are often criticized because they might encourage individuals to make less effort, which reduces the overall quality of the selected candidates. We show below that, in fact, the demographic parity mechanism always encourages the previously underrepresented group to make a larger effort than in the unconstrained case. For the previously overrepresented group, the situation varies: in many situations, the candidates from the overrepresented group will make a lower effort than in the unconstrained case, but when $\sqrt{C_L/C_H} < \alpha$ and $\alpha > p_L$, we show that they make a higher effort compared to that of the unconstrained case. Note that the last result may seem rather counterintuitive as demographic parity reduces the competition between groups. Later, in Section 5.3, we show the implications of this result on the average quality of the selected candidates.

COROLLARY 2 (EQUILIBRIUM EFFORT RATIO IN \mathcal{G}^{un} vs. \mathcal{G}^{dp}). *Let μ_H^{un} and μ_L^{un} be the equilibrium effort distributions in the unconstrained game \mathcal{G}^{un} . Similarly, let μ_H^{dp} and μ_L^{dp} be the equilibrium effort distributions in the game with the demographic parity constraint \mathcal{G}^{dp} .*

(i) *If $C_H = C_L$ and $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$, then*

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{un}}}{\bar{\mu}_H^{\text{dp}}} = \begin{cases} 1/p_H & \text{if } \alpha \leq p_H, \\ 1/\alpha & \text{if } \alpha > p_H, \end{cases} \quad \text{and} \quad \lim_{S \rightarrow \infty} \frac{\bar{\mu}_L^{\text{un}}}{\bar{\mu}_L^{\text{dp}}} = \begin{cases} 0 & \text{if } \alpha \leq p_H, \\ \frac{\alpha - p_H}{\alpha - \alpha p_H} & \text{if } \alpha > p_H. \end{cases}$$

(ii) If $C_H > C_L$, then

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{un}}}{\bar{\mu}_H^{\text{dp}}} = \begin{cases} 0 & \text{if } \alpha \leq p_L, \\ \frac{\alpha - p_L}{\alpha - \alpha p_L} & \text{if } \alpha > p_L, \end{cases} \quad \text{and} \quad \lim_{S \rightarrow \infty} \frac{\bar{\mu}_L^{\text{un}}}{\bar{\mu}_L^{\text{dp}}} = \begin{cases} 1/p_L & \text{if } \alpha \leq p_L, \\ \sqrt{\frac{C_L}{C_H}} \frac{1}{\alpha} & \text{if } \alpha > p_L. \end{cases}$$

PROOF SKETCH. Similarly, as in the proof of Corollary 1, we calculate the limits by using the equilibrium strategies found in Theorem 2 and in Theorem 4. The full proof is in Appendix A.6. \square

5.3 Selection quality with and without the demographic parity mechanism

We now show the implication of the previous result on how the demographic parity mechanism affects the selection quality. In a non-strategic setting (see e.g., [13]), the unconstrained decision maker is optimal in expectation. We show here that this no longer holds in the strategic setting: there exist scenarios under which *the Bayesian decision-maker is not optimal* in terms of the expected quality of selection. In such settings, the demographic parity constraint leads to a more qualified cohort. This counterintuitive phenomenon is due to the fact that demographic parity induces less competition between groups but a more fair competition within each group (compared to the unconstrained case).

We denote by Q^{un} and by Q^{dp} the expected quality at equilibrium of the selected candidates for the unconstrained and the demographic parity constrained games. In the theorem below, we characterize the ratio of the equilibrium cohort qualities by the unconstrained decision-maker from Section 4, and the demographic parity constrained decision-maker that we study in this section. For group-independent cost coefficient ($C_H = C_L$), we show that the quality ratio $Q^{\text{un}}/Q^{\text{dp}}$ tends to 1 in the limit of large S . For group-dependent cost coefficient ($C_H \neq C_L$), we show that the quality ratio can be smaller than one—the Bayesian decision-maker is not optimal if the candidates are strategic, and the demographic parity mechanism can lead to a better-qualified cohort.

THEOREM 5 (SELECTION QUALITY RATIO FOR \mathcal{G}^{un} AND \mathcal{G}^{dp}). *Let Q^{un} and Q^{dp} be the expected qualities of selection at equilibrium of the game \mathcal{G}^{un} and of the game \mathcal{G}^{dp} , respectively.*

- (i) *If $C_H = C_L$, then the ratio of the expected quality given by the unconstrained and the demographic parity constrained decision-makers tends to one with S : $\lim_{S \rightarrow \infty} \frac{Q^{\text{un}}}{Q^{\text{dp}}} = 1$.*
- (ii) *If $C_H > C_L$, then the ratio of the expected quality given by the unconstrained and the demographic parity constrained decision-makers can be smaller than one. Formally, for $c = \sqrt{C_L/C_H}$:*

$$\lim_{S \rightarrow \infty} \frac{Q^{\text{un}}}{Q^{\text{dp}}} = \begin{cases} \frac{1}{c p_H + p_L} > 1 & \text{if } p_L \geq \alpha, \\ \frac{c}{c p_H + p_L} < 1 & \text{if } p_L < \alpha. \end{cases}$$

PROOF SKETCH. First, we prove that as $S \rightarrow \infty$, the expected quality of the selected cohort grows at equal rate with the expected effort: $E(W_G \cdot [\tilde{W}_G \geq \theta]) \sim \bar{\mu}_G$. Hence, using the equilibrium strategies from Theorem 2 and Theorem 4, we can estimate the ratio $Q^{\text{un}}/Q^{\text{dp}}$ in the limit of $S \rightarrow \infty$. The complete proof is given in Appendix A.7. \square

We emphasize that the condition under which the demographic parity mechanism improves the average selection equality is when the selection rate α is larger than the size of the cost-advantaged group. The improvement of selection quality due to the demographic parity mechanism is explained by the fact that, without the demographic parity constraint, the advantaged minority has no incentives to make a large effort because the competition includes a lot of cost-disadvantaged candidates. Once the competition is among candidates of each separate populations, the cost-advantaged candidates have to compete with other cost-advantaged candidates, so they have to make a larger effort to be selected.

The demographic parity can decrease the average selection quality when $\alpha \leq p_L$ and when both groups have different cost coefficients C_G . In this case, if the low-noise L -candidates are the majority, then the ratio of quality $Q^{\text{un}}/Q^{\text{dp}}$ cannot be larger than 2 as S goes to infinity, regardless of the cost coefficients.

6 COMPLEMENTARY RESULTS

In the previous sections, we studied the properties of the Nash equilibrium of the game for large rewards S and showed that they can lead to discriminations. In this section, we complement this theoretical analysis by studying two (essentially independent) problems: First, does the population converge to the equilibrium? Second, what happens when the reward S is not large?

6.1 Convergence to the Nash equilibrium

To answer the first question, we perform a series of numerical experiments in which the decisions are made repeatedly. At a given time t , the candidates consider past data to make a strategic decision. This could represent, for instance, the case of college admission where candidates consider the distribution of grades from previous years; in this example, each decision epoch is a different year.

We study two population dynamics: *best response* and *fictitious play*.

- For the *best response* dynamics, at each of the discrete times $t \in \{1, 2, \dots, T\}$, the candidates observe the strategy played at the previous time step $\boldsymbol{\mu}^{(t-1)}$ and play a best response to it:

$$m_G^{(t)} \in b_G(\theta(\boldsymbol{\mu}^{(t-1)})).$$

- For the *fictitious play* dynamics, at each of the discrete times $t = 1, 2, \dots, T$, the candidates observe the whole history of plays and assume that the distribution of efforts is the empirical distribution of effort from time 1 to T . Candidates then play a best response to it:

$$m_G^{(t)} \in b_G(\theta(\hat{\boldsymbol{\mu}}^{(t)})),$$

$$\text{where } \theta(\hat{\boldsymbol{\mu}}^{(t)}) = F_{\hat{\boldsymbol{\mu}}^{(t)}}^{-1}(1 - \alpha) \text{ and } \hat{\boldsymbol{\mu}}^{(t)} = \sum_{s=1}^{t-1} \frac{1}{t-1} \boldsymbol{\mu}^{(s)}.$$

We numerically evaluate these two policies and report the results in Fig. 2. For the best response dynamics, we observe that $\mathbf{m}^{(t)} = (m_H^{(t)}, m_L^{(t)})$ converges to a limit cycle for any starting point. This is because when S is large,⁸ the best response map is not continuous (recall Fig. 1). The period of the limit cycle increases with the reward size S but the behavior is similar for all S : starting from $(0, 0)$, the candidates from both populations increase the effort as time increases. Then, the competition becomes too high and one of the populations drops out, i.e., make almost zero effort. After, the competition is only among the candidates of a single population until it becomes too difficult and all candidates drop out and return to the initial state. The cycle ends here, and the new cycle starts. In Fig. 2a and 2b, we also plot the average trajectory $\bar{\mathbf{m}}^{(t)} = (\bar{m}_H^{(t)}, \bar{m}_L^{(t)})$, where $\bar{m}_G^{(t)} = \frac{1}{t-1} \sum_{s=1}^{t-1} m_G^{(s)}$. We observe that the average over the trajectory seems to converge, yet the average effort over the trajectory is significantly larger than that of the average equilibrium effort for both groups.

The case of fictitious play dynamics is different, and it is depicted in Fig. 2c and 2d. In this case, the empirical distribution of efforts *does* converge to the Nash equilibrium. This is illustrated on the figure by the fact that the empirical average of effort converges to the average value of effort of the Nash equilibrium: $\bar{m}_G^{(t)} \xrightarrow{t \rightarrow \infty} \bar{\boldsymbol{\mu}}_G^{\text{un}}$ for both groups $G \in \{H, L\}$. Note that this is not a pointwise convergence but rather a convergence to a cycle: at equilibrium, the strategy played by the L -candidates converges to a cycle on the values b_L^{min} and b_L^{max} .

⁸For $S < \frac{1}{2} C_G \bar{\sigma}_G^2 / \phi(1)$ we can show that the function T is a contraction mapping, so any trajectory of the best response dynamics converges to the Nash equilibrium.




Fig. 2. Best response (BR) and fictitious play (FP) dynamics for different rewards S . The parameters of simulations are $T = 500$, $C_H = 1.5$, $C_L = 1$, $p_H = p_L = 0.5$, $\alpha = 0.1$, $\tilde{\sigma}_H = 0.6$ and $\tilde{\sigma}_L = 1$.

6.2 Selection problems with small and intermediate rewards

We start with the case of small rewards $S < C_G \tilde{\sigma}_G^2 / \phi(1)$ for which the results are quite different from the ones obtained for large S in Section 4. We show that the cost ratio C_L/C_H , as well as the expected quality variance ratio $\tilde{\sigma}_L^2/\tilde{\sigma}_H^2$, play significant roles in the outcome of the game. If the cost for the H -candidates is too high compared to that of the L -candidates, then the H -candidates make lower effort for all selection sizes α . Otherwise, if the cost for H -candidates is comparable to that of L -candidates (e.g., $C_H = C_L$) or it is lower, then for both small and large selection sizes α , the high-noise H -candidates make less effort compared to that of low-noise L -candidates. For both cases, there exists a parameter dependent value of α_0 such that for all values of $\alpha \leq \alpha_0$, the high-noise H -candidates are underrepresented. Interestingly, in (ii) of Theorem 6, we also observe that it is possible that H -candidates make a larger effort than L -candidates, yet they are selected at a lower rate, $\bar{x}_H^{\text{un}} < \bar{x}_L^{\text{un}}$.

Overall, the high-level interpretation of the below result is that *for small enough rewards S and for small enough values of α , the high-noise H -candidates are always underrepresented*. This result is in contrast with the results for large S studied in Section 4, and it is similar to the result in the non-strategic setting studied in [13, 16].

THEOREM 6 (DISCRIMINATION FOR SMALL REWARDS S). *Assume that $S < C_G \tilde{\sigma}_G^2 / \phi(1)$ for all $G \in \{H, L\}$. Let μ_H^{un} and μ_L^{un} , x_H^{un} and x_L^{un} be the equilibrium efforts and selection rates of the game \mathcal{G}^{un} .*

Denote $K_\mu := \sqrt{\frac{-2 \ln(C_H \tilde{\sigma}_H / (C_L \tilde{\sigma}_L))}{1/\tilde{\sigma}_H^2 - 1/\tilde{\sigma}_L^2}}$ and $K_x := \sqrt{\mathcal{W}\left(\frac{S^2 \left(\frac{1}{C_H \tilde{\sigma}_H} - \frac{1}{C_L \tilde{\sigma}_L}\right)^2}{2\pi(\tilde{\sigma}_L - \tilde{\sigma}_H)^2}\right)}$, where \mathcal{W} is the Lambert

function defined as the inverse to the function $f(\lambda) = \lambda e^\lambda$ and Φ^c is the complementary cumulative distribution function of the standard normal distribution $\mathcal{N}(0, 1)$.

- (i) If $C_H \tilde{\sigma}_H > C_L \tilde{\sigma}_L$, then $\bar{\mu}_H^{\text{un}} < \bar{\mu}_L^{\text{un}}$ for all $\alpha \in (0, 1)$, and $\bar{x}_H^{\text{un}} < \bar{x}_L^{\text{un}}$ if and only if $\alpha < \Phi^c(-K_x)$.
- (ii) If $C_H \tilde{\sigma}_H \leq C_L \tilde{\sigma}_L$, then

$$\bar{\mu}_H^{\text{un}} < \bar{\mu}_L^{\text{un}} \iff \alpha \in \left(0, \sum_{G \in \{A, B\}} p_G \Phi^c(K_\mu / \tilde{\sigma}_G)\right) \cup \left(\sum_{G \in \{A, B\}} p_G \Phi^c(-K_\mu / \tilde{\sigma}_G), 1\right),$$

$$\bar{x}_H^{\text{un}} < \bar{x}_L^{\text{un}} \iff \alpha < \Phi^c(K_x).$$

PROOF SKETCH. We show that for $S < C_G \tilde{\sigma}_G^2 / \phi(1)$ the best response in pure strategies is unique, hence the equilibrium effort distribution is a singleton $\mu_G^{\text{un}} = \delta(m - m_G^{\text{un}})$. Note that the first-order




Fig. 3. Characterization of the equilibrium for different rewards S . The parameters of simulations are $C_L = 1$, $p_L = 0.5$, $\tilde{\sigma}_H = 0.6$ and $\tilde{\sigma}_L = 1$.

condition on a maximum of the payoff function u_G is also a sufficient condition; it can be written:

$$\frac{S}{\tilde{\sigma}_G} \phi \left(\frac{m_G^{\text{un}} - \theta^{\text{un}}}{\tilde{\sigma}_G} \right) - C_G m_G^{\text{un}} = 0 \iff m_G^{\text{un}} = \frac{S}{C_G \tilde{\sigma}_G} \phi \left(\frac{m_G^{\text{un}} - \theta^{\text{un}}}{\tilde{\sigma}_G} \right).$$

Since we aim to find a value of α when $m_H^{\text{un}} = m_L^{\text{un}}$, we equate the right-hand sides of the above equation for two groups, and, by solving this equation, we obtain the values of $\theta^{\text{un}} - m_G^{\text{un}}$. By substituting this expression to the budget constraint, we derive the value of α at which $m_H^{\text{un}} = m_L^{\text{un}}$:

$$\alpha = \sum_{G \in \{H, L\}} p_G \Phi^c \left(\frac{\theta^{\text{un}} - m_G^{\text{un}}}{\tilde{\sigma}_G} \right).$$

The proof for x_G^{un} is similar to that of m_G^{un} . A complete proof is given in Appendix A.8. \square

Intermediate rewards. To conclude our analysis, we fill the gap between our theoretical results for the cases of small and large rewards S using numerical simulations.⁹ We perform our numerical simulations for the values of reward $S = 1, 10, 100, 1000$. The simulation result for $S = 1$ is studied theoretically in the first part of this section, as $S = 1$ satisfies the condition $S < C_G \tilde{\sigma}_G^2 / \phi(1)$. The result for $S = \infty$, studied in Section 4 and Section 5, is represented in Fig. 3 using a black solid line.

We plot the ratio of $\bar{x}_L^{\text{un}} / \bar{x}_H^{\text{un}}$ for the case of group-independent cost coefficient in Fig. 3a, the ratios of $\bar{x}_H^{\text{un}} / \bar{x}_L^{\text{un}}$ and $Q^{\text{un}} / Q^{\text{dp}}$ for the case of group-dependent cost coefficient in Fig. 3b and Fig. 3c. Overall, we observe a relatively smooth transition between the two regimes of S in all figures. In addition, the behavior for $S = 100$ and for $S = 1000$ is quite close to the behavior for $S = \infty$.

7 CONCLUDING DISCUSSION

In this work, we propose a simple model of selection with strategic candidates who are faced with group-dependent cost-of-effort and group-dependent noise variance. We characterize the resulting discrimination at equilibrium as well as the impact of removing it through the demographic parity mechanism that mandates equal representation across groups. Note that, in the context of our strategic model, demographic parity is not the only fairness notion that might make sense, and one could be tempted to consider, for instance, a meritocratic notion of fairness (where the representation would be commensurate with the effort). However, it remains important to understand the impact of imposing demographic parity as it is one of the most commonly used fairness notions.

⁹The code can be found at <https://gitlab.inria.fr/vemelian/strategic-selection-code>.

Throughout the paper, we made several simplifying assumptions, often to make the results easier to state and understand and to better isolate the effect of strategic behavior. Our work can be extended, however, in multiple ways:

Group-dependent variance of quality. We assumed that the variance of latent qualities is group-independent, i.e., $\eta_H^2 = \eta_L^2 = \eta^2$. This assumption can easily be removed, as all the results can equivalently be stated for $\tilde{\sigma}_H^2$ and $\tilde{\sigma}_L^2$ even with group-dependent variance. We make this assumption only for simplicity of exposition since it implies that $\tilde{\sigma}_H^2 < \tilde{\sigma}_L^2$ if and only if $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$.

Unobservable effort. In our model, we assume that the effort m is observable to the decision-maker. If the effort m is not observable, we argue that the decision-maker performs the selection based on the noisy estimate \hat{W} rather than the posterior expectation \bar{W} (which corresponds to the group-oblivious decision-maker in the terminology of [13]). All the results from this paper still hold but we need to replace the variance of the expected qualities $\tilde{\sigma}_G^2$ by the variance of the estimate $\eta^2 + \hat{\sigma}_G^2$. Note that most of the statements will be reversed as $\tilde{\sigma}_H^2 < \tilde{\sigma}_L^2$ if and only if $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$. In this case, the high-noise group has a higher variance of estimate, and the low-noise group has a lower variance of estimate. Hence, when the low-noise group is underrepresented in our model, the low-noise group is overrepresented in the model with unobservable effort.

More than two groups of candidates. In our model, we assume two groups of candidates, yet the results can be extended to more than two groups (e.g., in the proof of uniqueness of the equilibrium we do not rely on the fact that the number of populations is two). This additional dimension would simply add more interactions between the groups and complicate the statement of the results. For instance, in the case of multiple groups where two of them are subject to equal cost coefficients, we expect that the sorting will be with respect to the noise for these two groups, and with respect to the cost coefficients for the rest of the groups.

Monomial cost function and non-Gaussian noise. In our model, we assume a quadratic cost and Gaussian noise. We believe that these results will not change much if we assume other symmetric noise and monomial cost functions, i.e., $C_G m^{d_G}$. In this case, we expect that the best response could be characterized by a dropout threshold, as in our model. In addition, the power d_G of the monomial would be another important feature of the model: if it is group-dependent, then we expect that the dropout threshold will grow as $(S/C_G)^{1/d_G}$ and the candidates with higher d_G will drop out earlier, independently on the relations of C_G and $\tilde{\sigma}_G^2$.

Other models of candidate's utility. In our model, we assume rational risk-neutral candidates faced with different costs of efforts, and we also assume a risk-neutral Bayesian decision-maker. In non-strategic settings (see e.g., [13]), such a decision-maker is proven to be Bayesian-optimal, yet we prove that it is not optimal in the strategic setting. In our model, we do not consider other barriers for candidates as, for example, self-selection. To model self-selection, we can assume that there is a minimal threshold θ_G^{self} that each candidate should pass. The outcome of the equilibrium would then depend on the relation between θ_G^{self} and the dropout threshold θ_G^d . We can also easily consider a risk-averse (or risk loving) decision maker: in the case of an exponential utility function, this will lead to an additive bias f_G proportional to the variance $\tilde{\sigma}_G^2$, i.e., $\tilde{W}_i \sim \mathcal{N}(m_{G_i} + f_{G_i}, \tilde{\sigma}_{G_i}^2)$.

ACKNOWLEDGMENTS

This work has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003), by the French National Research Agency (ANR) through grant ANR-20-CE23-0007 and grant ANR-19-CE23-0015.

REFERENCES

- [1] Ali Ahmed, Mark Granberg, and Shantanu Khanna. 2021. Gender discrimination in hiring: An experimental reexamination of the Swedish case. *PLOS ONE* 16, 1 (01 2021), 1–15. <https://doi.org/10.1371/journal.pone.0245513>
- [2] Dennis J. Aigner and Glen G. Cain. 1977. Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review* 30, 2 (1977), 175–187.
- [3] Kenneth J. Arrow. 1973. The theory of discrimination. In *Discrimination in Labor Markets*, Orley Ashenfelter and Albert Rees (Eds.). Princeton University Press, 3–33.
- [4] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [5] Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review* 94, 4 (September 2004), 991–1013.
- [6] Mark Braverman and Sumegha Garg. 2020. The Role of Randomness and Noise in Strategic Classification. In *Proceedings of The Symposium on Foundations of Responsible Computing (FORC)*.
- [7] L. Elisa Celis, Chris Hays, Anay Mehrotra, and Nisheeth K. Vishnoi. 2021. The Effect of the Rooney Rule on Implicit Bias in the Long Term. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Mar 2021). <https://doi.org/10.1145/3442188.3445930>
- [8] L. Elisa Celis, Anay Mehrotra, and Nisheeth K. Vishnoi. 2020. Interventions for Ranking in the Presence of Implicit Bias. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT*)*, 369–380.
- [9] Stephen Coate and Glenn Loury. 1993. Will Affirmative-Action Policies Eliminate Negative Stereotypes? *American Economic Review* 83 (February 1993), 1220–40.
- [10] Patricia Cohen. 2019. New Evidence of Age Bias in Hiring, and a Push to Fight It. *The New York Times* (2019). <https://www.nytimes.com/2019/06/07/business/economy/age-discrimination-jobs-hiring.html>
- [11] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. 2018. Strategic Classification from Revealed Preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 55–70.
- [12] Vitalii Emelianov, Nicolas Gast, Krishna P. Gummadi, and Patrick Loiseau. 2020. On Fair Selection in the Presence of Implicit Variance. *Proceedings of the 21st ACM Conference on Economics and Computation* (Jul 2020). <https://doi.org/10.1145/3391403.3399482>
- [13] Vitalii Emelianov, Nicolas Gast, Krishna P. Gummadi, and Patrick Loiseau. 2022. On fair selection in the presence of implicit and differential variance. *Artificial Intelligence* 302 (2022), 103609. <https://doi.org/10.1016/j.artint.2021.103609>
- [14] Hanming Fang and Andrea Moro. 2011. Chapter 5 - Theories of Statistical Discrimination and Affirmative Action: A Survey. *Handbook of Social Economics*, Vol. 1. North-Holland, 133–200. <https://doi.org/10.1016/B978-0-444-53187-2.00005-X>
- [15] Qiang Fu and Zenan Wu. 2019. Contests: Theory and Topics. *Oxford Research Encyclopedia of Economics* (2019). <https://doi.org/10.1093/acrefore/9780190625979.013.440>
- [16] Nikhil Garg, Hannah Li, and Faidra Monachou. 2021. Standardized Tests and Affirmative Action: The Role of Bias and Variance. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 261. <https://doi.org/10.1145/3442188.3445889>
- [17] Jeannie Suk Gersen. 2019. The many sins of college admissions. *The New Yorker* (2019). <https://www.newyorker.com/news/our-columnists/the-many-sins-of-college-admissions>
- [18] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. 2016. Strategic Classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, 111–122.
- [19] Human Rights Watch. 2020. China: gender discrimination in hiring persists. *Human Rights Watch* (2020). <https://www.hrw.org/news/2020/04/29/china-gender-discrimination-hiring-persists>
- [20] Jon Kleinberg and Manish Raghavan. 2019. How Do Classifiers Induce Agents to Invest Effort Strategically?. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 825–844.
- [21] Jon M. Kleinberg and Manish Raghavan. 2018. Selection Problems in the Presence of Implicit Bias. In *Proceedings of the 9th Innovations in Theoretical Computer Science Conference (ITCS)*, 33:1–33:17.
- [22] Edward Lazear and Sherwin Rosen. 1981. Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy* 89, 5 (1981), 841–64. <https://EconPapers.repec.org/RePEc:ucp:jpolec:v:89:y:1981:i:5:p:841-64>
- [23] Lydia T. Liu, Nikhil Garg, and Christian Borgs. 2022. Strategic ranking. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research)*, Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (Eds.), Vol. 151. PMLR, 2489–2518. <https://proceedings.mlr.press/v151/liu22b.html>
- [24] Shelly J. Lundberg and Richard Startz. 1983. Private Discrimination and Social Intervention in Competitive Labor Market. *The American Economic Review* 73, 3 (1983), 340–347. <http://www.jstor.org/stable/1808117>
- [25] Anay Mehrotra, Bary S. R. Pradeliski, and Nisheeth K. Vishnoi. 2022. Selection in the Presence of Implicit Bias: The Advantage of Intersectional Constraints. arXiv:cs.CY/2202.01661

- [26] John Miller, Smitha Milli, and Moritz Hardt. 2020. Strategic Classification is Causal Modeling in Disguise. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
- [27] Smitha Milli, John Miller, Anca D. Dragan, and Moritz Hardt. 2019. The Social Cost of Strategic Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*. 230–239.
- [28] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. arXiv:cs.IR/2201.12662
- [29] Edmund Phelps. 1972. The Statistical Theory of Racism and Sexism. *American Economic Review* 62, 4 (1972), 659–61.
- [30] Lincoln Quillian, Devah Pager, Arnfinn H. Midtbøen, and Ole Hexel. 2017. Hiring Discrimination Against Black Americans Hasn’t Declined in 25 Years. *Harvard Business Review* (2017). <https://hbr.org/2017/10/hiring-discrimination-against-black-americans-hasnt-declined-in-25-years>
- [31] William H Sandholm. 2010. *Population games and evolutionary dynamics*. MIT press.
- [32] Andrew Schotter and Keith Weigelt. 1992. Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results. *The Quarterly Journal of Economics* 107, 2 (05 1992), 511–539. <https://doi.org/10.2307/2118480>
- [33] Stratis Tsirtsis and Manuel Gomez-Rodriguez. 2020. Decisions, Counterfactual Explanations and Strategic Behavior. In *Proceedings of the Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS)*.
- [34] Suzanne Woolley. 2017. How More Americans Are Getting a Perfect Credit Score. *Bloomberg* (2017). <https://www.bloomberg.com/news/features/2017-08-14/obsessives-have-cracked-the-perfect-fico-credit-score-of-850>
- [35] Hanrui Zhang, Yu Cheng, and Vincent Conitzer. 2019. When Samples Are Strategically Selected. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. 7345–7353.

A MISSING PROOFS

A.1 Properties of the best response

Recall that the payoff function of an individual with effort m and given the selection threshold θ is

$$u(m; \theta) = S \cdot \Phi\left(\frac{m - \theta}{\tilde{\sigma}}\right) - \frac{1}{2}Cm^2,$$

where Φ is the CDF of the standard normal distribution.

Denoting ϕ the PDF of the standard normal distribution, the first two derivatives of u with respect to m are:

$$\begin{aligned}\frac{\partial u}{\partial m}(m; \theta) &= \frac{S}{\tilde{\sigma}}\phi\left(\frac{\theta - m}{\tilde{\sigma}}\right) - Cm, \\ \frac{\partial^2 u}{(\partial m)^2}(m; \theta) &= \frac{S}{\tilde{\sigma}}\phi\left(\frac{\theta - m}{\tilde{\sigma}}\right)\frac{\theta - m}{\tilde{\sigma}^2} - C.\end{aligned}$$

The payoff function u is defined on $[0, \infty)$, and it is continuous and continuously differentiable. Moreover, $\frac{\partial u}{\partial m}(m = 0) > 0$ and $u(m; \theta) \xrightarrow{m \rightarrow \infty} -\infty$. Hence, all local maxima of u must satisfy the first-order condition (FOC) $\frac{\partial u}{\partial m}(m; \theta) = 0$ and the second-order condition (SOC) $\frac{\partial^2 u}{\partial m^2} \leq 0$. The maximum of the payoff function is attained in one of the local maxima.

We start by a first lemma.

LEMMA 3 (MAXIMA OF u). *Fix S , C and $\tilde{\sigma}$.*

- (i) *If $S < C\tilde{\sigma}^2/\phi(1)$, then there is a unique global maximum of $u(m; \theta)$ for all θ .*
- (ii) *If $S \geq C\tilde{\sigma}^2/\phi(1)$, then there exists a unique $\theta^d(S)$ such that for $\theta = \theta^d(S)$ there are two global maxima of $u(m; \theta)$, and for $\theta \neq \theta^d$, there is a unique global maximum of $u(m; \theta)$.*


PROOF. Let us denote by $z = (m - \theta)/\tilde{\sigma}$ and let $v(z) = \frac{S}{\tilde{\sigma}}\phi(z) - C\tilde{\sigma}z$ and $w(z) := -\frac{S}{\tilde{\sigma}}z\phi(z) - C\tilde{\sigma} = dv(z)/dz$. The first and second derivatives of u can be expressed as a function of v and w :

$$\begin{aligned}\frac{\partial u}{\partial m}(m; \theta) &= v(z) - C\theta, \\ \frac{\partial^2 u}{(\partial m)^2}(m; \theta) &= \frac{1}{\tilde{\sigma}}w(z).\end{aligned}$$

The function $z \cdot \phi(z)$ has the global maximum at $z = 1$ which is equal to $\phi(1)$, and the global minimum at $z = -1$ which is equal to $-\phi(1)$. Hence, two cases are possible:

- (i) If $\phi(1) < C\tilde{\sigma}^2/S$, then $w(z) < 0$ for all $z \neq 1$. Hence, v is a strictly decreasing function (since $w < 0$ in this case), so the FOC gives a unique solution m which is a global maximum of u .
- (ii) If $\phi(1) \geq C\tilde{\sigma}^2/S$, then the equation $w(z) = 0$ has two real solutions, denoted by $z_1 \leq z_2$. They can be found explicitly, i.e., $z_{1,2} = -\sqrt{-\mathcal{W}_{1,2}\left(-\frac{2\pi C^2 \tilde{\sigma}^4}{S^2}\right)}$ where \mathcal{W} is the Lambert function defined as the inverse to the function $f(y) = ye^y$. Note also that $z_1 \xrightarrow{S \rightarrow \infty} -\infty$ and $z_2 \xrightarrow{S \rightarrow \infty} 0$.

We consider the latter case (ii) in details. We can verify, that the function v is a decreasing function for $z \in (-\infty, z_1) \cup (z_2, \infty)$, and it is an increasing function for $z \in (z_1, z_2)$. This shows that the function $v(z)$ has the same shape as the curve on the right. As a result, the FOC condition $v(z) - C\theta = 0$ can have at most three solutions depending on the value of θ .



Indeed, let $\theta_1 = v(z_1)/C$ and $\theta_2 = v(z_2)/C$.

- For all $\theta \notin (\theta_1, \theta_2)$, the FOC gives a unique solution which is a global maximum of $u(m; \theta)$.
- For all $\theta \in (\theta_1, \theta_2)$, we have that $v(z_1) \leq C\theta$ and $v(z_2) \geq C\theta$ which guarantees three solutions to FOC which we denote by $m_1 \leq m_2 \leq m_3$.
- For $\theta = \theta_1$, we have $u(m_1(\theta); \theta) = u(m_2(\theta); \theta)$, and for $\theta = \theta_2$ we have that $u(m_2(\theta); \theta) = u(m_3(\theta); \theta)$. For all $\theta \in (\theta_1, \theta_2)$, the m_2 is a local minimum, m_1 and m_3 are local maxima.

Consider the two continuous, differentiable and non-negative functions $\Delta_{12}(\theta) = u(m_1(\theta); \theta) - u(m_2(\theta); \theta)$ and $\Delta_{32}(\theta) = u(m_3(\theta); \theta) - u(m_2(\theta); \theta)$. We can see that $\Delta_{12}(\theta_1) = 0$ and $\Delta_{32}(\theta_2) = 0$.

We can also verify that $\frac{\partial u(m(\theta), \theta)}{\partial \theta} = Cm\left(\frac{\partial m}{\partial \theta} - 1\right) - Cm\frac{\partial m}{\partial \theta} = -Cm(\theta) < 0$. Therefore, the function Δ_{12} is increasing, since $\frac{d\Delta_{12}}{d\theta} = C(m_2(\theta) - m_1(\theta)) > 0$, whereas the function Δ_{32} is decreasing since $\frac{d\Delta_{32}}{d\theta} = C(m_2(\theta) - m_3(\theta)) < 0$. Hence, there is a unique $\theta^d \in (\theta_1, \theta_2)$, such that $\Delta_{12}(\theta^d) = \Delta_{32}(\theta^d)$ which is equivalent to $u(m_1(\theta^d), \theta^d) = u(m_3(\theta^d), \theta^d)$. □

A.2 Proof of Lemma 1

We start with the proof of the case (i). According to Lemma 3, there are two pure best response values, b^{\max} and b^{\min} , that correspond to the dropout threshold θ^d . Following the definition of the expected payoff $u(m; \theta)$, we can show the following upper bound on the values of pure best responses $b(\theta^d) \in \{b^{\max}(\theta^d), b^{\min}(\theta^d)\}$:

$$0 < u(b(\theta^d); \theta^d) = S \cdot x(b(\theta^d); \theta^d) - \frac{Cb^2(\theta^d)}{2} \leq S - \frac{Cb^2(\theta^d)}{2}.$$

This implies that $b(\theta^d) \leq \sqrt{2S/C}$.

Second, we show that $|\theta^d - b(\theta^d)|$ is not bounded as $S \rightarrow \infty$. By assuming that $|\theta^d - b(\theta^d)| < \varepsilon$ for all S , we end up with the following contradictory inequality that must hold for any value of S :

$$\phi\left(\frac{\varepsilon}{\tilde{\sigma}}\right) < \phi\left(\frac{\theta^d - b(\theta^d)}{\tilde{\sigma}}\right) = \frac{C\tilde{\sigma}}{S} \cdot b(\theta^d) \leq \tilde{\sigma}\sqrt{2C/S} \xrightarrow{S \rightarrow \infty} 0,$$

where ϕ is the PDF of the standard normal distribution $\mathcal{N}(0, 1)$.

Since the value of $|\theta^d - b(\theta^d)|$ is not bounded, by studying the first and the second derivatives of the payoff function (as in Lemma 3), we can show that $(b^{\max}(\theta^d) - \theta^d) \xrightarrow{S \rightarrow \infty} +\infty$ and $(b^{\min}(\theta^d) - \theta^d) \xrightarrow{S \rightarrow \infty} -\infty$. Hence, the selection rates, corresponding to the b^{\max} and b^{\min} converge to:

$$\begin{aligned} \lim_{S \rightarrow \infty} x(b^{\max}(\theta^d); \theta^d) &= \lim_{S \rightarrow \infty} \Phi\left(\frac{b^{\max}(\theta^d) - \theta^d}{\tilde{\sigma}}\right) = 1, \\ \lim_{S \rightarrow \infty} x(b^{\min}(\theta^d); \theta^d) &= \lim_{S \rightarrow \infty} \Phi\left(\frac{b^{\min}(\theta^d) - \theta^d}{\tilde{\sigma}}\right) = 0, \end{aligned}$$

where Φ is the CDF of the standard normal distribution $\mathcal{N}(0, 1)$.

Asymptotic behavior of b^{\min} and b^{\max} . Using the definition of the dropout threshold θ^d and the definition for b^{\min} , b^{\max} , we can write:

$$\begin{aligned} 0 < u(b^{\max}(\theta^d); \theta^d) &= u(b^{\min}(\theta^d); \theta^d) \iff \\ 0 < x(b^{\max}(\theta^d); \theta^d) - \frac{C(b^{\max}(\theta^d))^2}{2S} &= x(b^{\min}(\theta^d); \theta^d) - \frac{C(b^{\min}(\theta^d))^2}{2S} < x(b^{\min}(\theta^d); \theta^d). \end{aligned}$$

As $\lim_{S \rightarrow \infty} x(b^{\min}(\theta^d); \theta^d) = 0$ and $\lim_{S \rightarrow \infty} x(b^{\max}(\theta^d); \theta^d) = 1$, it implies that

$$\lim_{S \rightarrow \infty} \frac{C(b^{\max}(\theta^d))^2}{2S} = 1 \text{ and } \lim_{S \rightarrow \infty} \frac{C(b^{\min}(\theta^d))^2}{2S} = 0.$$

Hence, we show the asymptotic behavior of the pure strategy best response at θ^d for $S \rightarrow \infty$:

$$b^{\max}(\theta^d(S)) = \sqrt{2S/C}(1 + o(1)) \text{ and } b^{\min}(\theta^d(S)) = o(\sqrt{2S/C}).$$

Asymptotic behavior of θ^d . Finally, we consider the asymptotic behavior of the dropout threshold θ^d . The dropout threshold $\theta^d(S)$ is unbounded: if we assume the opposite, then b^{\max} would be $o(1)$, $S \rightarrow \infty$ since b^{\max} must satisfy the FOC:

$$b^{\max}(\theta^d) = \frac{S}{C\tilde{\sigma}} \phi\left(\frac{\theta^d - b^{\max}(\theta^d)}{\tilde{\sigma}}\right).$$

Hence, consider the following limit which we find using l'Hôpital rule and the properties of b^{\min} and b^{\max} proven above:

$$\lim_{S \rightarrow \infty} \frac{\theta^d(S)}{\sqrt{2S/C}} = \lim_{S \rightarrow \infty} \frac{\frac{b^{\max}(\theta^d(S)) + b^{\min}(\theta^d(S))}{2S}}{\frac{1}{2}S^{-1/2}\sqrt{2/C}} = \lim_{S \rightarrow \infty} \frac{b^{\max}(\theta^d(S)) + b^{\min}(\theta^d(S))}{\sqrt{2S/C}} = 1.$$

Refined asymptotic behavior of b^{\min} . Since $b^{\min}(\theta^d(S)) = o(\sqrt{2S/C})$ and $\theta^d(S) = \sqrt{2S/C}(1 + o(1))$ as $S \rightarrow \infty$, then, using the first-order condition, we show

$$b^{\min}(\theta^d(S)) = \frac{S}{C\tilde{\sigma}} \phi\left(\frac{\theta^d(S) - b^{\min}(\theta^d(S))}{\tilde{\sigma}}\right) = \frac{S}{C\tilde{\sigma}} \phi\left(\frac{\sqrt{2S/C}(1 + o(1))}{\tilde{\sigma}}\right) \xrightarrow{S \rightarrow \infty} 0.$$

Now, we are ready to proof the case (ii). According to Lemma 3, the pure best response is unique. Consider two cases:

- If $\theta(S) = \gamma\theta^d(S)$, then $\lim_{S \rightarrow \infty} (b(\gamma\theta^d(S)) - \gamma\theta^d(S)) = +\infty$. Consider the following limit that we calculate using l'Hôpital rule:

$$\lim_{S \rightarrow \infty} \frac{b(\gamma\theta^d(S))}{\gamma\theta^d(S)} = \lim_{S \rightarrow \infty} \frac{db/d\theta \cdot d\theta/dS}{d\theta/dS} = \lim_{S \rightarrow \infty} \frac{db}{d\theta} = \lim_{S \rightarrow \infty} \frac{b(b-\theta)}{b(b-\theta) + \tilde{\sigma}^2} = 1.$$

- If $\theta(S) = \theta^d(S)/\gamma$, then $\lim_{S \rightarrow \infty} (b(\theta^d(S)/\gamma) - \theta^d(S)/\gamma) = -\infty$. Since the payoff function u is non-negative, we have:

$$u \geq 0 \iff Cb(\theta^d(S)/\gamma)^2/(2S) \leq \Phi\left(\frac{b(\theta^d(S)/\gamma) - \theta^d(S)/\gamma}{\tilde{\sigma}}\right) \xrightarrow{S \rightarrow \infty} 0.$$

Hence, $b(\theta^d(S)/\gamma) = o(\sqrt{2S/C})$. Given that the best response must satisfy the FOC:

$$b(\theta^d(S)/\gamma) = \frac{S}{C\tilde{\sigma}} \phi\left(\frac{\theta^d(S)/\gamma - b(\theta^d(S)/\gamma)}{\tilde{\sigma}}\right) = \frac{S}{C\tilde{\sigma}} \phi\left(\frac{\sqrt{2S/C} \frac{1}{\gamma}(1 + o(1))}{\tilde{\sigma}}\right) = o(1).$$

A.3 Proof of Lemma 2

Case (i). The asymptotic behavior of the dropout threshold is given in Appendix A.2. We recall the proof here. Consider the following limit which we find using l'Hôpital rule and the properties of b^{\min} and b^{\max} proven in Lemma 1:

$$\lim_{S \rightarrow \infty} \frac{\theta^d(S)}{\sqrt{2S/C}} = \lim_{S \rightarrow \infty} \frac{\frac{b^{\max}(\theta^d(S)) + b^{\min}(\theta^d(S))}{2S}}{\frac{1}{2}S^{-1/2}\sqrt{2/C}} = \lim_{S \rightarrow \infty} \frac{b^{\max}(\theta^d(S)) + b^{\min}(\theta^d(S))}{\sqrt{2S/C}} = 1.$$

Case (ii). To prove that the dropout threshold $\theta^d(S; \tilde{\sigma})$ is decreasing with $\tilde{\sigma}$, we differentiate with respect to $\tilde{\sigma}$ the condition on the equal payoffs for b^{\min} and b^{\max} strategies, and we obtain:

$$\frac{\partial \theta^d}{\partial \tilde{\sigma}} = -\frac{b^{\max}(\theta^d(S)) + b^{\min}(\theta^d(S)) - \theta^d(S)}{\tilde{\sigma}} < 0 \text{ for large enough } S.$$

A.4 Proof of Theorem 1

Let us denote by T the set-valued function

$$T(\theta) = \{F_{\mu}^{-1}(1 - \alpha) : \mu_G \in \beta_G(\theta)\},$$

where $\mu = (\mu_H, \mu_L)$ and F_{μ} is the CDF of \tilde{W} induced by μ : $F_{\mu} = p_H F_{\mu_H} + p_L F_{\mu_L}$.

LEMMA 4. *There is a one-to-one correspondence between the fixed points of T and the equilibria μ^{un} of the game \mathcal{G}^{un} .*

PROOF. If μ^{un} is an equilibrium of the game \mathcal{G}^{un} , then there is a unique θ^{un} such that $F_{\mu^{\text{un}}}^{-1}(1 - \alpha) = \theta^{\text{un}}$ since $F_{\mu^{\text{un}}}(\theta)$ is a monotone function. This θ^{un} is a fixed point of T since $\mu_G^{\text{un}} \in \beta_G(\theta^{\text{un}})$ by definition of μ^{un} .

On the other hand, if θ^{un} is a fixed point of T , then there is a unique $\mu^{\text{un}} = (\mu_H^{\text{un}}, \mu_L^{\text{un}})$ such that $\mu_G^{\text{un}} \in \beta_G(\theta^{\text{un}})$:

- (1) If $\theta \neq \theta_G^d$ for all $G \in \{H, L\}$, then the pure best responses b_H and b_L are unique (see Lemma 3 in Appendix A.1), so μ^{un} is unique.

(2) If, without loss of generality, $\theta = \theta_H^d$, then there is a unique pure best response b_L , and two pure response values b_H^{\max}, b_H^{\min} (see Lemma 3 in Appendix A.1). The function $F(\tau) = F_{\mu}(\theta)$ is monotone in $\tau \in [0, 1]$ which is a probability for $\mu_H = \tau \delta(m - b_H^{\max}) + (1 - \tau) \delta(m - b_H^{\min})$. It is also an equilibrium of \mathcal{G}^{un} as $\theta = F_{\mu}^{-1}(1 - \alpha)$ by definition of θ . Hence, $\mu^{\text{un}} = (\mu_H, \mu_L)$ is an equilibrium of the game \mathcal{G}^{un} and it is unique. \square

By showing a one-to-one correspondence between equilibria of \mathcal{G}^{un} and the fixed points of T , it is left us to prove that the function $T(\theta)$ has a unique fixed point θ^{un} , i.e., a solution to $T(\theta^{\text{un}}) = \theta^{\text{un}}$ is unique. This will imply that the equilibrium of the game \mathcal{G}^{un} is unique. We use the following lemmas.

LEMMA 5. Assume that $m(\theta)$ satisfies FOC and SOC defined in Section A.1.

- (i) If $m > \theta$, then $0 < \frac{dm}{d\theta} < 1$.
- (ii) If $m \leq \theta$, then $\frac{dm}{d\theta} \leq 0$.

PROOF. First, we derive the expression for the first derivative. We differentiate the both sides of FOC defined in Appendix A.1:

$$\frac{dm}{d\theta} = \frac{S}{C\tilde{\sigma}} \left(\frac{m - \theta}{\tilde{\sigma}} \right) \phi \left(\frac{\theta - m}{\tilde{\sigma}} \right) \left(\frac{1 - \frac{dm}{d\theta}}{\tilde{\sigma}} \right) = \frac{m - \theta}{\tilde{\sigma}^2} m \left(1 - \frac{dm}{d\theta} \right) \implies \frac{dm}{d\theta} = \frac{m(m - \theta)}{m(m - \theta) + \tilde{\sigma}^2}.$$

Second, from the SOC defined in Appendix A.1, we find that $m(m - \theta) + \tilde{\sigma}^2 > 0$. Hence, the sign of the derivative $dm/d\theta$ is determined only by the sign of its numerator $m(m - \theta)$. Note that the value of m is strictly positive as it satisfies the FOC, so $dm/d\theta > 0$ if and only if $m > \theta$. \square

LEMMA 6. For all S and $\theta \neq \theta_G^d(S)$, the total derivative $dT/d\theta$ can be expressed as:

$$\frac{dT}{d\theta} = \frac{\sum_G p_G \frac{1}{\tilde{\sigma}_G} \phi \left(\frac{T - b_G}{\tilde{\sigma}_G} \right) \cdot \frac{b_G(b_G - \theta)}{b_G(b_G - \theta) + \tilde{\sigma}_G^2}}{\sum_G p_G \frac{1}{\tilde{\sigma}_G} \phi \left(\frac{T - b_G}{\tilde{\sigma}_G} \right)} < 1.$$

PROOF. Since T is an implicit function of θ , we use the chain rule to find the total derivative:

$$\frac{dT}{d\theta} = \sum_G \frac{\partial T}{\partial b_G} \frac{db_G}{d\theta}.$$

By differentiating both sides of the budget constraint, we can find the partial derivative $\partial T/\partial m_G$:

$$\frac{\partial}{\partial b_G} \left(\sum_G p_G x_G(b_G; T) \right) = 0 \iff \frac{\partial T}{\partial b_G} = \frac{p_G \frac{1}{\tilde{\sigma}_G} \phi \left(\frac{T - b_G}{\tilde{\sigma}_G} \right)}{\sum_G p_G \frac{1}{\tilde{\sigma}_G} \phi \left(\frac{T - b_G}{\tilde{\sigma}_G} \right)}.$$

Finally, using the fact that $\partial T/\partial b_G > 0$, where $\sum_G \partial T/\partial b_G = 1$ and Lemma 5, we show that $dT/d\theta < 1$. \square

Existence. First, let us show that there is an interval $[\theta_0, \theta_1]$, such that for all $\theta \in \mathbf{R}$, we have that $\theta_0 \leq T(\theta) \leq \theta_1$:

- Since the best response $b_G(\theta) \geq 0$ for all $G \in \{H, L\}$, then for any fixed α , let θ_0 be the solution to the equation $\sum_G p_G \Phi^c \left(\frac{\theta_0 - 0}{\tilde{\sigma}_G} \right) = \alpha$. Hence, $T(\theta) \geq \theta_0$ for all $\theta \in \mathbf{R}$.

- Since the best response $b_G(\theta) \leq \sqrt{2S/C_G}$ for all $G \in \{H, L\}$, then for any fixed α , let θ_1 be the solution to the equation $\sum_G p_G \Phi^c \left(\frac{\theta_1 - \sqrt{2S/C_G}}{\sigma_G} \right) = \alpha$. Hence, $T(\theta) \leq \theta_1$ for all $\theta \in \mathbb{R}$.

Therefore, we can consider the function T on the interval $[\theta_0, \theta_1]$ which is compact and convex. The graph of the function T is closed and for all $\theta \in [\theta_0, \theta_1]$, we have that $T(\theta)$ is convex (for $\theta \neq \theta_G^d$, the value of $T(\theta)$ is unique, and for $\theta = \theta_G^d$, the value of $T(\theta)$ is an interval). Hence, using Kakutani fixed point theorem, we show that the fixed point exists.

Uniqueness. We show that there exists a fixed point of the function T . For all $\theta \neq \theta_G^d$, we have that $dT/d\theta < 1$ (Lemma 6), and we have that $\lim_{\theta \rightarrow \theta_G^d+} T(\theta) \leq \lim_{\theta \rightarrow \theta_G^d-} T(\theta)$ for all $G \in \{H, L\}$. Hence, there is a unique solution to the fixed-point problem $T(\theta) = \theta$, and, as a result, a unique equilibrium of the game \mathcal{G}^{un} due to Lemma 4.

A.5 Proof of Theorem 2

Case (i). We prove that the dropout threshold $\theta_1^d(S)$ is the fixed point of T , i.e., it corresponds to the equilibrium of the game \mathcal{G}^{un} . As we show in the proof of Lemma 1, as $S \rightarrow \infty$, we have $x_G^{\max} := x(b^{\max}(\theta_G^d); \theta_G^d) \xrightarrow{S \rightarrow \infty} 1$ and $x_G^{\min} := x(b^{\min}(\theta_G^d); \theta_G^d) \xrightarrow{S \rightarrow \infty} 0$.

If $\alpha \leq p_1$, and given that the selection rate for the G_2 -candidates at θ_1^d tends to zero with S , assume that G_1 -candidates randomize their strategy by playing b_1^{\max} with the probability τ_1 and b_1^{\min} with the probability $1 - \tau_1$. The G_2 -candidates play $b_2(\theta_1^d) \xrightarrow{S \rightarrow \infty} 0$. The probability τ_1 can be found from the budget constraint: $\alpha = p_1(\tau_1 x_1^{\max} + (1 - \tau_1) x_1^{\min}) + p_1 x_2(\theta_1^d, b_2(\theta_1^d)) \xrightarrow{S \rightarrow \infty} p_1 \tau_1$, so $\tau_1 \xrightarrow{S \rightarrow \infty} \alpha/p_1$. This strategy satisfies the budget constraint, so θ_1^d is the fixed point of T .

Case (ii). We now prove that the dropout threshold θ_2^d is fixed point of T . If $\alpha > p_1$, then selecting all candidates from G_1 group would not be enough, and some G_2 -candidates are needed to satisfy the selection rate α .

Given θ_2^d , the fraction of selected G_1 -candidates would be equal to $x_1(\theta_2^d, b_1(\theta_2^d)) \xrightarrow{S \rightarrow \infty} 1$. We claim the G_2 -candidates will play b_2^{\max} with probability τ_2 , and b_2^{\min} with probability $1 - \tau_2$. The probability τ_2 can be found from the budget constraint:

$$\alpha = p_1 x_1(\theta_2^d, b_1(\theta_2^d)) + p_2 (x_2^{\max} \tau_2 + x_2^{\min} (1 - \tau_2)) \xrightarrow{S \rightarrow \infty} p_1 + p_2 \tau_2.$$

Hence, for $\tau_2 \xrightarrow{S \rightarrow \infty} \frac{\alpha - p_1}{p_2}$, the dropout threshold, θ_2^d is the fixed point of T .

A.6 Proof of Corollary 2

We use the expressions for equilibrium strategies from Theorem 2 and Theorem 4.

- (i) If $C_H = C_L = C$ and $\hat{\sigma}_H^2 > \hat{\sigma}_L^2$:

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{un}}}{\bar{\mu}_H^{\text{dp}}} = \lim_{S \rightarrow \infty} \begin{cases} \frac{\alpha/p_H \sqrt{2S/C}(1+o(1))}{\alpha \sqrt{2S/C}(1+o(1))} & \text{if } \alpha \leq p_H \\ \frac{\sqrt{2S/C}(1+o(1))}{\alpha \sqrt{2S/C}(1+o(1))} & \text{if } \alpha > p_H \end{cases} = \begin{cases} 1/p_H & \text{if } \alpha \leq p_H, \\ 1/\alpha & \text{if } \alpha > p_H, \end{cases}$$

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_L^{\text{un}}}{\bar{\mu}_L^{\text{dp}}} = \lim_{S \rightarrow \infty} \begin{cases} \frac{o(1)}{\alpha \sqrt{2S/C}(1+o(1))} & \text{if } \alpha \leq p_H \\ \frac{\alpha - p_H}{p_L} \frac{\sqrt{2S/C}(1+o(1))}{\alpha \sqrt{2S/C}(1+o(1))} & \text{if } \alpha > p_H \end{cases} = \begin{cases} 0 & \text{if } \alpha \leq p_H, \\ \frac{\alpha - p_H}{\alpha - \alpha p_H} & \text{if } \alpha > p_H. \end{cases}$$

(ii) If $C_H > C_L$:

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_H^{\text{un}}}{\bar{\mu}_H^{\text{dp}}} = \lim_{S \rightarrow \infty} \begin{cases} \frac{o(1)}{\alpha \sqrt{2S/C_H(1+o(1))}} & \text{if } \alpha \leq p_L \\ \frac{(\alpha - p_L)/p_H \sqrt{2S/C_H(1+o(1))}}{\alpha \sqrt{2S/C_H(1+o(1))}} & \text{if } \alpha > p_L \end{cases} = \begin{cases} 0 & \text{if } \alpha \leq p_L, \\ \frac{\alpha - p_L}{\alpha - \alpha p_L} & \text{if } \alpha > p_L, \end{cases}$$

$$\lim_{S \rightarrow \infty} \frac{\bar{\mu}_L^{\text{un}}}{\bar{\mu}_L^{\text{dp}}} = \lim_{S \rightarrow \infty} \begin{cases} \frac{\alpha/p_L \sqrt{2S/C_L(1+o(1))}}{\alpha \sqrt{2S/C_L(1+o(1))}} & \text{if } \alpha \leq p_L \\ \frac{\sqrt{2S/C_H(1+o(1))}}{\alpha \sqrt{2S/C_L(1+o(1))}} & \text{if } \alpha > p_L \end{cases} = \begin{cases} 1/p_L & \text{if } \alpha \leq p_L, \\ \sqrt{\frac{C_L}{C_H}} \frac{1}{\alpha} & \text{if } \alpha > p_L. \end{cases}$$

A.7 Proof of Theorem 5

First, using the formula for the expected value of the truncated normal distribution, we find:

$$\begin{aligned} Q^{\text{un}}(m_G; \theta) &= \sum_G p_G \mathbb{E}(W_G \cdot [\tilde{W}_G > \theta_G]) = \sum_G p_G \mathbb{E}(W_G | \tilde{W}_G > \theta_G) \mathbb{P}(\tilde{W}_G \geq \theta_G) \\ &= \sum_G p_G \left[m_G \Phi\left(\frac{\theta_G - m_G}{\tilde{\sigma}_G}\right) + \tilde{\sigma}_G \phi\left(\frac{\theta_G - m_G}{\tilde{\sigma}_G}\right) \right]. \end{aligned}$$

Second, using Lemma 1, we can show that for all $\gamma \in (0, 1)$ and $\theta_G = \gamma \cdot \theta_G^d$, we have:

$$b_G(\theta_G) \cdot \Phi\left(\frac{\theta_G - b_G(\theta_G)}{\tilde{\sigma}_G}\right) + \tilde{\sigma}_G \phi\left(\frac{\theta_G - b_G(\theta_G)}{\tilde{\sigma}_G}\right) = \theta_G(1 + o(1))(1 + o(1)) + \tilde{\sigma}_G \cdot o(1) = \theta_G(1 + o(1))$$

and, for $\theta_G = \theta_G^d/\gamma$, we have:

$$b_G(\theta_G) \cdot \Phi\left(\frac{\theta_G - b_G(\theta_G)}{\tilde{\sigma}_G}\right) + \tilde{\sigma}_G \phi\left(\frac{\theta_G - b_G(\theta_G)}{\tilde{\sigma}_G}\right) \xrightarrow{S \rightarrow \infty} 0.$$

For $\theta_G = \theta_G^d$, we have:

$$\begin{aligned} b_G^{\max}(\theta_G) \cdot \Phi\left(\frac{\theta_G - b_G^{\max}(\theta_G)}{\tilde{\sigma}_G}\right) + \tilde{\sigma}_G \phi\left(\frac{\theta_G - b_G^{\max}(\theta_G)}{\tilde{\sigma}_G}\right) &= \theta_G(1 + o(1)), \\ b_G^{\min}(\theta_G) \cdot \Phi\left(\frac{\theta_G - b_G^{\min}(\theta_G)}{\tilde{\sigma}_G}\right) + \tilde{\sigma}_G \phi\left(\frac{\theta_G - b_G^{\min}(\theta_G)}{\tilde{\sigma}_G}\right) &\xrightarrow{S \rightarrow \infty} 0. \end{aligned}$$

For the game \mathcal{G}^{dp} , we use the equilibrium strategy from Theorem 4, and can write that:

$$\begin{aligned} Q^{\text{dp}} &= p_H(\alpha + o(1)) \cdot \theta_H^d(S)(1 + o(1)) + p_L \cdot (\alpha + o(1)) \theta_L^d(S)(1 + o(1)) \\ &= \alpha p_H \cdot \theta_H^d(S)(1 + o(1)) + \alpha p_L \cdot \theta_L^d(S)(1 + o(1)), S \rightarrow \infty. \end{aligned}$$

For the game \mathcal{G}^{un} , we use the equilibrium strategy from Theorem 2. We consider separately the case of group-independent and the case of group-dependent costs.

Group-independent cost.

If $\alpha \leq p_H$, then

$$Q^{\text{un}} = p_H(\alpha/p_H + o(1)) \cdot \theta_H^d(1 + o(1)) + p_L \cdot o(1) = \alpha \cdot \sqrt{2S/C}(1 + o(1)), S \rightarrow \infty.$$

If $\alpha > p_H$, then

$$Q^{\text{un}} = p_H(1 + o(1)) \cdot \theta_L^d(1 + o(1)) + p_L \cdot \left(\frac{\alpha - p_H}{p_L} + o(1)\right) \theta_L^d(1 + o(1)) = \alpha \sqrt{2S/C}(1 + o(1)), S \rightarrow \infty,$$

where we use the assumption on equal costs $C_H = C_L$ and the result of Lemma 2 on the asymptotic behavior of the dropout threshold. Hence, we can calculate the limit:

$$\lim_{S \rightarrow \infty} Q^{\text{dp}}/Q^{\text{un}} = 1.$$

Group-dependent cost.

If $\alpha \leq p_L$, then

$$Q^{\text{un}} = p_L(\alpha/p_L + o(1)) \cdot \theta_L^{\text{d}} \cdot (1 + o(1)) + p_H \cdot o(1) = \alpha \cdot \sqrt{2S/C_L} \cdot (1 + o(1)).$$

If $\alpha > p_L$, then

$$\begin{aligned} Q^{\text{un}} &= p_L(1 + o(1)) \cdot \theta_H^{\text{d}} \cdot (1 + o(1)) + p_H \cdot \left(\frac{\alpha - p_L}{p_H} + o(1) \right) \theta_H^{\text{d}}(1 + o(1)) \\ &= p_L(1 + o(1))\theta_H^{\text{d}} \cdot (1 + o(1)) + (\alpha - p_L)\theta_H^{\text{d}}(1 + o(1)). \end{aligned}$$

Hence, for $c := \sqrt{C_L/C_H}$, we can write that:

$$\lim_{S \rightarrow \infty} Q^{\text{un}}/Q^{\text{dp}} = \begin{cases} \lim_{S \rightarrow \infty} \frac{\alpha \theta_L^{\text{d}}(1+o(1))}{\alpha p_H \theta_H^{\text{d}}(1+o(1)) + \alpha p_L \theta_L^{\text{d}}(1+o(1))} = \frac{1}{c p_H + p_L} & \text{if } \alpha \leq p_L, \\ \lim_{S \rightarrow \infty} \frac{p_L \theta_H^{\text{d}}(1+o(1)) + (\alpha - p_L) \theta_H^{\text{d}}(1+o(1))}{\alpha p_H \theta_H^{\text{d}}(1+o(1)) + \alpha p_L \theta_L^{\text{d}}(1+o(1))} = \frac{c}{c p_H + p_L} & \text{if } \alpha > p_L, \end{cases}$$

where $\frac{c}{c p_H + p_L} < \frac{c}{c p_H + c p_L} = 1$.

A.8 Proof of Theorem 6

Equilibrium efforts. Let us first find for which $\alpha \in (0, 1)$ we have the equality of effort at equilibrium, i.e., $m_H^{\text{un}} = m_L^{\text{un}}$. Since the equilibrium values are the best responses to θ^{un} , and the pure best response is unique (see Lemma 3), we can write:

$$\begin{aligned} m_H^{\text{un}} = m_L^{\text{un}} &\iff \frac{S}{C_H \tilde{\sigma}_H} \phi\left(\frac{\theta^{\text{un}} - m_H^{\text{un}}}{\tilde{\sigma}_H}\right) = \frac{S}{C_L \tilde{\sigma}_L} \phi\left(\frac{\theta^{\text{un}} - m_L^{\text{un}}}{\tilde{\sigma}_L}\right) \\ &\iff \theta^{\text{un}} - m_H^{\text{un}} = \pm \sqrt{-2 \frac{\log(C_H \tilde{\sigma}_H) - \log(C_L \tilde{\sigma}_L)}{1/\tilde{\sigma}_H^2 - 1/\tilde{\sigma}_L^2}}. \end{aligned}$$

For $C_H \tilde{\sigma}_H > C_L \tilde{\sigma}_L$, there exist no real solution to the above equation. In this case, we can show that $m_H^{\text{un}} < m_L^{\text{un}}$ for all $\alpha \in (0, 1)$. Assume the opposite, i.e., $m_H^{\text{un}} > m_L^{\text{un}}$ for all θ , then $\phi((\theta - m_H^{\text{un}})/\tilde{\sigma}_H) < \phi((\theta - m_L^{\text{un}})/\tilde{\sigma}_L)$ and also $C_H \tilde{\sigma}_H > C_L \tilde{\sigma}_L$ which contradicts the initial assumption.

For $C_H \tilde{\sigma}_H \leq C_L \tilde{\sigma}_L$, there are two values of $\alpha_{m_H^{\text{un}}=m_L^{\text{un}}}$ which correspond to equal equilibrium efforts $m_H^{\text{un}} = m_L^{\text{un}}$:

$$\alpha_{m_H^{\text{un}}=m_L^{\text{un}}}^{(1,2)} = \sum_G p_G \Phi^c \left(\pm \frac{\sqrt{-2 \frac{\log(C_H \tilde{\sigma}_H) - \log(C_L \tilde{\sigma}_L)}{1/\tilde{\sigma}_H^2 - 1/\tilde{\sigma}_L^2}}}{\tilde{\sigma}_G} \right).$$

We can verify that $db_H/d\alpha > db_L/d\alpha$ for $\alpha_{m_H^{\text{un}}=m_L^{\text{un}}}^{(1)}$ and $db_H/d\alpha < db_L/d\alpha$ for $\alpha_{m_H^{\text{un}}=m_L^{\text{un}}}^{(2)}$ which concludes the proof.

Equilibrium selection rates. Let us find such $\alpha_{x_H=x_L}$ for which both groups are selected at equal rates, i.e., $x_H^{\text{un}} = x_L^{\text{un}}$. This is equivalent to:

$$\Phi^c\left(\frac{\theta^{\text{un}} - m_H^{\text{un}}}{\tilde{\sigma}_H}\right) = \Phi^c\left(\frac{\theta^{\text{un}} - m_L^{\text{un}}}{\tilde{\sigma}_L}\right) = \alpha_{x_H=x_L} \iff \theta^{\text{un}} = m_G^{\text{un}} + \tilde{\sigma}_G \Phi^{-1}(1 - \alpha_{x_H=x_L}), \quad \forall G \in \{H, L\}.$$

By definition of the best response, we can also write that:

$$m_G^{\text{un}} = \frac{S}{C_G \tilde{\sigma}_G} \phi\left(\frac{\theta^{\text{un}} - m_G^{\text{un}}}{\tilde{\sigma}_G}\right) = \frac{S}{C_G \tilde{\sigma}_G} \phi\left(\Phi^{-1}(1 - \alpha_{x_H=x_L})\right),$$

$$\theta^{\text{un}} = \frac{S}{C_G \tilde{\sigma}_G} \phi\left(\Phi^{-1}(1 - \alpha_{x_H=x_L})\right) + \tilde{\sigma}_G \Phi^{-1}(1 - \alpha_{x_H=x_L}).$$

The equality for θ^{un} is possible only for such α , when

$$\frac{S}{C_H \tilde{\sigma}_H} \phi\left(\Phi^{-1}(1 - \alpha_{x_H=x_L})\right) + \tilde{\sigma}_H \Phi^{-1}(1 - \alpha_{x_H=x_L}) = \frac{S}{C_L \tilde{\sigma}_L} \phi\left(\Phi^{-1}(1 - \alpha_{x_H=x_L})\right) + \tilde{\sigma}_L \Phi^{-1}(1 - \alpha_{x_H=x_L}),$$

which is if and only if

$$S\left(\frac{1}{C_H \tilde{\sigma}_H} - \frac{1}{C_L \tilde{\sigma}_L}\right) \phi\left(\Phi^{-1}(1 - \alpha_{x_H=x_L})\right) = \Phi^{-1}(1 - \alpha_{x_H=x_L})(\tilde{\sigma}_L - \tilde{\sigma}_H).$$

We solve the corresponding equation by letting $z := \Phi^{-1}(1 - \alpha_{x_H=x_L})$ and solving with respect to z . After, $\alpha_{x_H=x_L} = 1 - \Phi(z)$.

Let $\xi = S\left(\frac{1}{C_H \tilde{\sigma}_H} - \frac{1}{C_L \tilde{\sigma}_L}\right) / (\tilde{\sigma}_L - \tilde{\sigma}_H)$. Then by definition of ϕ :

$$\xi \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} = z \iff z \cdot e^{\frac{z^2}{2}} = \xi \frac{1}{\sqrt{2\pi}}.$$

By squaring both sides of the above equation, we end up with the equation of type $y \cdot e^y = \text{const}$ which can be solved using the Lambert \mathcal{W} function defined as the inverse to $f(y) = y \cdot e^y$. As a result, by solving for z^2 and taking the square root with the sign equal to the sign of ξ , we show:

$$\alpha_{x_H=x_L} = 1 - \Phi\left(\text{sgn}(\xi) \sqrt{\mathcal{W}\left(\frac{\xi^2}{2\pi}\right)}\right).$$

We can verify that the $dx_H/d\alpha > dx_L/d\alpha$ at $\alpha_{x_H=x_L}$, so for all $\alpha < \alpha_{x_H=x_L} = 1 - \Phi\left(\text{sgn}(\xi) \sqrt{\mathcal{W}\left(\frac{\xi^2}{2\pi}\right)}\right)$, we have that $x_H^{\text{un}} < x_L^{\text{un}}$.