Algorithmic Bias and the Diversity Paradox in Leadership

Finance Working Paper N° 1104/2025
November 2025

Pierre Chaigneau
Queen's University and ECGI

ECGI Working Paper Series in Finance

Algorithmic Bias and the Diversity Paradox in Leadership

Working Paper N° 1104/2025
November 2025

Pierre Chaigneau

Abstract

We study how algorithmic biases in recruitment affect top management diversity. Individuals differ in talent and invest in human capital before appointment. When appointment thresholds differ systematically across groups, a higher threshold lowers entry-level representation but can raise representation at the top. By motivating the most talented individuals to invest more, higher thresholds generate a thicker upper tail of the ability distribution. Consequently, diversity at senior levels may increase even as entry-level diversity declines. The model challenges the ``pipeline argument'' and predicts that algorithmic biases may improve top-management diversity in meritocratic firms but reduce it in less meritocratic ones.

Keywords: algorithmic bias, algorithmic fairness, discrimination, human capital, pipeline argument

JEL Classifications: G34, J24, J71, M51

Pierre Chaigneau
Associate Professor of Finance & Commerce '77 Fellow of Finance
Queen's University
143 Union Street West
Kingston, ON K7L 2P3, Canada
e-mail: pierre.chaigneau@queensu.ca

Algorithmic Bias and the Diversity Paradox in Leadership

Pierre Chaigneau
Smith School of Business at Queen's University

October 31, 2025

Abstract

We study how algorithmic biases in recruitment affect top management diversity. Individuals differ in talent and invest in human capital before appointment. When appointment thresholds differ systematically across groups, a higher threshold lowers entry-level representation but can raise representation at the top. By motivating the most talented individuals to invest more, higher thresholds generate a thicker upper tail of the ability distribution. Consequently, diversity at senior levels may increase even as entry-level diversity declines. The model challenges the "pipeline argument" and predicts that algorithmic biases may improve top-management diversity in meritocratic firms but reduce it in less meritocratic ones.

1 Introduction

The representation of demographic groups within firms has received substantial attention and motivated initiatives aimed at correcting historical imbalances. Some of these imbalances can be attributed to biased appointment processes (Goldin and Rouse (2000), Bertrand and Duflo (2017)). In response, many firms have modified appointment criteria to account for individuals' demographic characteristics. A large scale study found that two-thirds of US companies have strategies for hiring underrepresented groups, and 70% track gender representation in hiring.[1] Thus, the differential treatment of different groups in the appointment process is systematic, i.e. it is common to a large set of firms.

In turn, greater hiring of certain groups in entry-level positions can become self-reinforcing when recruitment processes rely on algorithms or artificial intelligence (AI), which are being adopted with increasing frequency (Köchling and Wehner (2020), Fabris et al. (2025)). Indeed, if an algorithm or AI learns that some type of candidate was more likely to be successfully hired in the past, it will tend to favor this type of candidate in the future. Moreover, it can be viewed as a black box: while the resulting outcomes are observable, the internal logic that generates them is typically opaque.[2] As a result, the differential appointment thresholds implicitly applied to different groups are effectively systematic but exogenous.

This has implications for top management representation. In principle, increasing the proportion of employees from underrepresented groups should ultimately improve the representation of these groups in top management positions. Indeed, companies often blame an insufficient "pipeline" of suitably qualified candidates to justify a lack of senior level diversity.[3] For example, a shortage of suitable candidates explains the rarity of female CEOs (He and Whited (2023)). In this perspective, it is surprising that the increase in historically underrepresented groups' representation in entry-level positions has not led to similar increases in upper-level representation (Fernandez-Mateo and Fernandez (2016)).[4] Larcker and Tayan (2020) note the lack of "tangible progress in increasing the prevalence of diverse executives in corporate leadership

[1]Source: Women in the Workplace 2017, by McKinsey and Lean In.

[2]Proprietary recruitment algorithms often rely on complex models trained on large datasets, making it difficult to trace how candidate characteristics affect appointment probabilities (Raghavan et al. (2020), Cowgill and Tucker (2020)). Moreover, organizations rarely disclose the structure or training data of these algorithms, both for intellectual property reasons and due to the inherent complexity of modern machine-learning systems (Köchling and Wehner (2020)).

[3]"One reason companies say they find it difficult to appoint more women to their boards is the lack of a "pipeline" of suitably qualified candidates." Source: *Financial Times*, Women still mind the gap on pay and power, July 14 2015.

[4]For example, in the US, women hold 47% of entry-level positions, but they account for only 37% of managers, 29% of Vice Presidents, 21% of C-suite leaders, and only 6% of Fortune 500 CEOs – a proportion that has been stable for years. Source: Women in the Workplace 2017, by McKinsey and Lean In. The survey included 222 companies employing more than 12 million people. Likewise, the *Financial Times* notes that "black employees have only seen significant growth in the most junior roles in finance" with a decline in senior role representation from 2007 to 2018. Source: *Financial Times*, Share of black employees in senior US finance roles falls despite diversity efforts, March 31 2021.

positions."

This paper develops a simple theoretical model to study how group-contingent appointment thresholds affect group representation in junior and senior positions. Unlike Coate and Loury (1993), individuals differ in talent, invest in human capital ("effort") to improve their ability, and are appointed to entry-level positions based on noisy estimates of their ability. The assumption of heterogeneous innate talent results in heterogeneous effort incentives and leads to novel predictions. We show that a lower appointment threshold for a particular group can increase group representation at the entry level while reducing it at the senior level. This challenges the common "pipeline argument" (Helfat et al., 2006), which would predict that greater representation at the junior level automatically translates into improved representation at the top. The model thus provides a potential explanation for observed patterns in which certain groups are relatively well represented at the entry level but remain underrepresented in top management, even absent initial differences between groups.

We assume that a candidate will only be appointed to an entry-level position if her estimated ability exceeds a given "appointment threshold". In our model, there can be a different appointment threshold for each group. We consider systematic appointment biases, including algorithmic biases, over which firms have little control. We study their implications for the probability that the firm CEO will belong to a particular group. To this end, we simply assume that the firm will select its CEO among the subset of appointed individuals with a minimum level of ability. This minimum level of ability required is interpreted as the level of "meritocracy" in the firm. In our model, different firms can be more or less meritocratic.

As a first step, we study the implications of these appointment thresholds for human capital accumulation decisions optimally made by individuals. In our model, individuals are atomistic, so that their effort decisions depend on the appointment threshold to entry-level positions, not on the probability of being appointed CEO since it is equal to zero. For example, undergraduate students who implicitly make an effort decision when choosing their courses and their majors will consider the effects of these choices on the probability of getting a job at a prestigious firm, but they will not consider how they change the probability of being appointed CEO of that firm.

A candidate will be especially motivated to exert "effort" by independently acquiring additional skills if she believes that enhancing her talent will have a sufficiently important effect on the appointment decision. At one end of the spectrum, individuals who are very talented will view such effort as superfluous, because they expect their estimated ability to be above the threshold regardless. At the other end of the spectrum, untalented individuals will think that effort is useless, because they expect their estimated ability to be below the threshold regardless. Thus, only individuals who believe that their talent is close to the appointment threshold will be motivated to enhance it by exerting effort. This mechanism is consistent with the empirical evidence in Huesmann et al. (2025).

Consequently, as in Coate and Loury (1993) and the large literature that followed, raising the threshold can either increase or decrease human capital investment by an individual.[5] In our model, raising the threshold increases the probability that a given individual will exert effort if and only if the threshold is below a certain level. Furthermore, this level is increasing in the individual's talent: talented individuals are best motivated by high appointment thresholds, i.e., when it is hard to be appointed to entry-level positions.

We use this framework to study the implications of group-contingent appointment thresholds for entry-level appointments and CEO selection.

First of all, we consider the case when information is perfect and symmetric: the ability of any individual is common knowledge, as in Krishna and Tarasov (2016). Our first main result is that, in any given group, the proportion of appointed individuals with ability high enough to be CEO is increasing in the appointment threshold for this group – as long as the threshold is lower than the minimum level of ability required to be CEO. Thus, in highly meritocratic firms, i.e. firms where this minimum level is high, the CEO is most likely to be from a group with a relatively high appointment threshold, even though the measure of appointed individuals from this group will be relatively low.

Intuitively, in a group that faces a relatively high appointment threshold, only talented individuals will exert effort to improve their ability. These individuals will then be very able, since they are initially talented and they also improved their human capital. By contrast, in a group that faces a relatively low appointment threshold, only untalented individuals will exert effort to improve their ability. Many individuals from this group will be appointed, but most will not be very able.

Next, we consider the case when individuals receive noisy signals about their talent before making their effort decisions, and are appointed based on their imperfectly measured ability.

First, we derive numerically the measure of appointed individuals from a group who are eligible for the CEO position as a function of the appointment threshold for this group, for several values of the minimum level of ability necessary to be eligible. We find that this measure is non-monotonic in the appointment threshold, and reaches its maximum for an interior value of the appointment threshold. Intuitively, only talented individuals have a chance to be selected as CEO, and setting the threshold either too low or too high discourages them from exerting effort.

Second, we compute numerically the appointment threshold that maximizes this measure. We find that this threshold is an increasing function of the level of meritocracy. Intuitively, highly meritocratic firms will only consider individuals with very high ability for the CEO position. Very few individuals from any group will have the minimum ability required based on

[5] Coate and Loury (1993) show that changing the threshold affects individual's effort decisions. In their model, however, individuals do not have heterogeneous "talent": only their effort decision ("investment") matters for employability purposes. This assumption is widely used in the literature that followed (e.g. Moro and Norman (2003, 2004), Fryer and Loury (2013)).

talent alone. Moreover, when individuals are (imperfectly) informed about their own talent, a high appointment threshold tends to elicit effort among highly talented individuals.

Third, we assume that the population is equally divided into two groups with the same talent distribution, and we compute numerically the probability that the firm CEO will be from one group or the other. As long as appointment thresholds are not too high, the CEO is most likely to be from the group with the highest appointment threshold, which will also be under-represented in entry-level jobs. The more meritocratic the firm, the larger the subset of appointment thresholds such that this result holds. On the contrary, in firms which are not very meritocratic, this tradeoff between group representation at the junior and senior level may not arise. Then, a bias against a group may simply reduce its representation both at the junior and senior levels.

Our main prediction is that applying different appointment thresholds to different groups will lead to opposite effects on group representation at the junior and senior level when the selection of senior management is highly meritocratic. This is in contrast to a naïve interpretation of the pipeline argument, which would predict that group representation at the senior level is facilitated by enhancing group representation at more junior levels.

Well-governed large firms are likely to be meritocratic. When CEOs have a multiplicative effect on firm value (Gabaix and Landier (2008)), even a small reduction in managerial ability can destroy billions of dollars in value for a large firm. Of course, this does not imply that all large firms are meritocratic. First, those with weak governance and severe agency problems may still fail to act in shareholders' interests (Kaplan and Minton (2011), Pagano and Picariello (2025)). Second, family firms may prioritize control over profitability (Burkart, Panunzi, and Shleifer (2003)). Overall, large non-family firms with effective governance mechanisms tend to operate meritocratically, although there are significant differences across countries (Bandiera et al., 2024).

In summary, when appointment thresholds remain moderate, the model predicts that large, well-governed firms will be more likely to be led by individuals from groups that face mild but systematic discrimination in recruitment algorithms. In contrast, it predicts that less meritocratic or poorly governed firms exhibit the opposite pattern. The analysis therefore reveals a potential unintended consequence of efforts to eliminate algorithmic bias.

2 Related literature

This paper contributes to a large literature on management diversity (Adams and Ferreira (2009), Nielsen (2010), Matsa and Miller (2011), Neely et al. (2020)) with a focus on top management appointments. For brevity and simplicity, it abstracts from several important factors that affects diversity in top management appointments (Smith et al. (2013), Cook and Glass (2014), Dezsö et al. (2016)). It is also not the first to highlight limitations of the pipeline argument (Castleman and Allen (1998), Cannady et al. (2014)), although to our knowledge it is the first to highlight the limitation based on human capital investments that we study. This paper is also related to the large economics literature on diversity policies (Fang and Moro (2011), Fryer and Loury (2013), Krishna and Tarasov (2016), Schaede and Mankki (2022)), including the literature that focuses on the private sector (Coate and Loury (1993), Peck (2017)). Finally, it is related to the large literature on human capital accumulation, especially in a context when groups are treated differently (Cotton et al. (2025)). Sethi and Somanathan (2023) and Chowdhury, Esteve-González, and Mukherjee (2023) review this literature.

Our theoretical model builds on a literature that follows Coate and Loury (1993) and which is surveyed by Fang and Moro (2011). Coate and Loury (1993) and many papers that followed, including Fryer (2007), do not assume heterogeneous talent. In these models, changes in the threshold do not have different implications for individuals with different talent – which is crucial in our analysis. Moreover, as opposed to Coate and Loury (1993), Fryer (2007), and other models of "statistical discrimination", the mechanism at play in our model does not involve multiple equilibria based on self-fulfilling beliefs that result in statistical discrimination.

We now discuss other closely related papers. Krishna and Tarasov (2016) focus on the propensity of individuals with different ability to exert effort to pass an admission exam, and study excessive effort due to the "rat race." As in our model, innate talent is identically distributed across groups; however, unlike our model, individuals in different groups vary in their capacity to improve their ability. In a model of search, Fershtman and Pavan (2021) assume that the evaluation of minority candidates is noisier than that of non-minorities, and show that this can make soft affirmative action policies counterproductive. Blair and Chung (2021) study the value of a signal of ability ("occupational licensing") in a model with workers with heterogeneous ability when the cost of obtaining this signal varies across groups. Wages are determined by firms that compete for workers and use observable licensing to update their beliefs about a worker's ability, similar to Spence (1973). Sethi and Somanathan (2018) study a setting where individuals with heterogeneous ability can enhance their ability through observable effort ("training"). When different groups do not have the same access to training opportunities, they show that threshold policies are not necessarily meritocratic.

In the models of Pikulina and Ferreira (2025) and Ferreira, Nikolowa, and Pikulina (2025), intrinsically biased employers can discriminate between equally qualified candidates. In these models, a candidate can be either skilled or unskilled, whereas two different candidates will not be equally qualified in our model. Pikulina and Ferreira (2025) study the consequences for human capital accumulation decisions across groups, and show that the bias can lead to either overinvestment or underinvestment depending on the stakes. Ferreira, Nikolowa, and Pikulina (2025) study the consequences for equilibrium outcomes, including labor market segmentation, firm profitability and economic efficiency.

3 The model

Each individual belongs to an identifiable group, denoted by k. An individual i has "talent" a_i.[6] It is drawn from the distribution $\tilde{a}_i \sim \mathcal{N}(\mu_k, \sigma_k^2)$. Thus, the parameters μ_k and σ_k^2 are respectively the mean and the variance of talent within group k. We assume that all groups have the same average talent, and we normalize $\mu_k = 0\ \forall k$. An individual privately receives an imperfectly informative signal s_i of her talent a_i: given talent a_i, the signal is $\tilde{s}_i = a_i + \tilde{\epsilon}_i^s$, where $\tilde{\epsilon}_i^s \sim \mathcal{N}(0, \sigma_s^2)$. Let s_i refer to the realization of the signal observed by individual i.

After observing signal s_i, each individual i can exert unobservable effort e_i to improve her ability. As in Coate and Loury (1993) and Krishna and Tarasov (2016), this effort captures human capital accumulation decisions privately made by individuals such as acquiring additional skills. Individual i chooses $e_i \in \{0, \overline{e}\}$ at cost $C(e_i)$, with $C(0) = 0$ and $C(\overline{e}) = c \in (0, 1)$. We henceforth distinguish between an individual's talent, a_i, and ability $a_i + e_i$.[7]

Appointments are based on noisy nonverifiable signals of ability. The signal on individual i's ability is $\tilde{q}_i = a_i + e_i + \tilde{\epsilon}_i^q$, where $\tilde{\epsilon}_i^q \sim \mathcal{N}(0, \sigma_q^2)$. An individual i from group k is appointed if and only if $q_i \geq t_k$. This assumption is standard following the seminal paper of Coate and Loury (1993). Thus, the appointment threshold can be contingent on the individual's group, and the threshold for group k is denoted by t_k. We are studying the impact of exogenous *systematic* biases, including algorithmic biases (Chen (2023)), on effort and appointment decisions.[8]

Individual i derives private benefits B from being appointed, which we normalize at $B = 1$, and chooses effort e_i to maximize:[9]

$$\mathbb{P}\left(\tilde{q}_i \geq t_k | s_i\right) - C(e_i) \tag{1}$$

The firm's CEO will be selected among all appointed individuals with ability greater than a target level $\bar{a}$: even though CEO appointments involve many different factors outside the scope of this paper, only high-ability individuals are considered for this position. This is consistent with the notion that a firm can only get noisy signals about the ability of outsiders (hence noisy appointments to entry-level positions), but can learn the ability of its employees with high potential. Indeed, the empirical evidence suggests that large firms mostly rely on insiders for CEO appointments because of information frictions (Cziraki and Jenter (2022)).

We normalize the measure of candidates to a continuum $[0, 1]$. The timeline of the model

[6] We use the standard definition of "talent" as natural aptitude or skill.

[7] Thus, talent and effort are perfect substitutes. In the model of Coate and Loury (1993), by contrast, an individual is "qualified" only if she has exerted effort.

[8] Algorithmic decision-making in hiring is often a black box: while the resulting outcomes are observable, the underlying decision rules are typically opaque and proprietary (Raghavan et al. (2020), Cowgill and Tucker (2020)). Accordingly, we treat group-dependent appointment thresholds as systematic but exogenous (Köchling and Wehner (2020), Fabris et al. (2025)). It is also noteworthy that idiosyncratic biases, which are specific to a single organization, would not have the same effect on individual effort decisions.

[9] As we make clear later, only the ratio c/B matters for effort decisions.

for each individual i from group k is as follows:

a_i drawn $\boxed{\text{variance } \sigma_k^2}$	individual i observes s_i $\boxed{\text{variance } \sigma_s^2}$	individual i makes effort decision	entry-level appointment if $q_i \geq t_k$ $\boxed{\text{variance } \sigma_q^2}$	CEO selected among high ability employees

In the model of Coate and Loury (1993), there is no heterogeneity in "talent" among individuals ($\sigma_k = 0$). In the model of Krishna and Tarasov (2016), individuals know their talent ($\sigma_s = 0$) and individual performance reveals their ability ($\sigma_q = 0$). Our results in section 5.1 involve a similar setting. In section 5.2, we also allow for individual uncertainty about one's talent ($\sigma_s > 0$), and imperfect screening processes ($\sigma_q > 0$).

4 Effort decision and appointments

4.1 Effort decision

We study the effort decision for a given appointment threshold. With a continuum of candidates, the probability for any individual to be appointed CEO is zero. Therefore, an individual does not consider a potential CEO appointment when she chooses her effort. Individuals update their beliefs after observing a signal of their talent. Individual i from group k's updated belief about her talent is normally distributed with the following mean and variance:

$$\mu_{a_i|s_i} \equiv \mathbb{E}\left[\tilde{a}_i|s_i\right] = \mu_k + \frac{cov(\tilde{a}_i, \tilde{s}_i)}{var(\tilde{s}_i)}\left[s_i - \mathbb{E}[\tilde{s}_i]\right] = \mu_k + \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right] \tag{2}$$

$$\sigma^2_{a_i|s_i} \equiv var(\tilde{a}_i|s_i) = \sigma_k^2 - \frac{\left(cov(\tilde{a}_i, \tilde{s}_i)\right)^2}{var(\tilde{s}_i)} = \frac{\sigma_k^2\sigma_s^2}{\sigma_k^2 + \sigma_s^2} \tag{3}$$

Conditional on s_i and e_i, the individual's performance $\tilde{q}_i$ is normally distributed with mean $e_i + \mu_k + \frac{\sigma_k^2}{\sigma_k^2+\sigma_s^2}\left[s_i - \mu_k\right]$, and variance $\Sigma^2 \equiv \frac{\sigma_k^2\sigma_s^2}{\sigma_k^2+\sigma_s^2} + \sigma_q^2$. Denoting by Φ the standard normal CDF, let s^* be the solution to:

$$\Phi\left(\Sigma^{-1}\left(t_k - \mu_k - \tfrac{\sigma_k^2}{\sigma_k^2+\sigma_s^2}\left[s^* - \mu_k\right]\right)\right) - \Phi\left(\Sigma^{-1}\left(t_k - \overline{e} - \mu_k - \tfrac{\sigma_k^2}{\sigma_k^2+\sigma_s^2}\left[s^* - \mu_k\right]\right)\right) = c. \tag{4}$$

For any given t_k, we assume that c is sufficiently low for equation (4) to have two solutions s^* (otherwise any individual would always exert zero effort). We denote these solutions by $\underline{s}(t_k)$ and $\overline{s}(t_k)$, with $\underline{s}(t_k) < \overline{s}(t_k)$, and we refer to them as the "effort bounds."

Lemma 1 *An individual i exerts effort if and only if her private signal $s_i \in [\underline{s}(t_k), \overline{s}(t_k)]$, and the sensitivity of these effort bounds to t_k is equal to $\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}$.*

Lemma 1 shows that the decision to exert effort is nonmonotonic in the signal that each individual receives about her talent. An individual who receives a very high signal ($s_i > \overline{s}$) is very confident about her ability to exceed the threshold, even without effort, while an individual who receives a very low signal ($s_i < \overline{s}$) believes that she will not exceed the threshold even if she exerts effort, so she does not. On the contrary, for individuals who receive signals within the bounds ($s_i \in [\underline{s}, \overline{s}]$), exerting effort has a sufficiently important effect on the likelihood to exceed the threshold, i.e., returns to effort are high, and these individuals optimally exert effort. Lemma 1 also shows how moving the threshold changes effort decisions: the effort bounds are linearly increasing in the appointment threshold t_k.

4.2 Appointments

We now study the appointment of individuals from group k given the appointment threshold t_k.

We take into account individuals' effort decisions, as derived in Lemma 1. Letting ϕ be the PDF of the signal distribution conditional on talent, the probability that an individual with talent a_i exceeds the threshold and gets appointed is:

$$\begin{aligned}
\mathbb{P}\left(\tilde{q}_i \geq t_k | a_i\right) &= \int_{-\infty}^{\infty} \mathbb{P}\left(a_i + e_i + \tilde{\epsilon}_i^q \geq t_k | a_i, s_i\right) \phi(s_i|a_i) ds_i \\
&= \int_{-\infty}^{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\right) \phi(s_i|a_i) ds_i \\
&\quad + \int_{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i - \overline{e}\right)\right)\right) \phi(s_i|a_i) ds_i \\
&\quad + \int_{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\infty} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\right) \phi(s_i|a_i) ds_i \qquad (5)
\end{aligned}$$

An increase in the threshold has two effects on the probability of appointment in equation (5). The first is a direct effect, which represents the change in the probability that performance exceeds the threshold for a given effort. This direct effect is always negative: an increase in the threshold lowers the probability that performance exceeds the threshold, ceteris paribus. The second is an incentive effect, which represents the change in the probability that performance exceeds the threshold due to a change in the effort decision. Formally:

$$\frac{d}{dt_k}\mathbb{P}\left(\tilde{q}_i \geq t_k | a_i\right) = \underbrace{\int_{-\infty}^{\infty} \frac{\partial}{\partial t_k}\mathbb{P}\left(a_i + e_i + \tilde{\epsilon}_i^q \geq t_k | a_i, s_i\right)\phi(s_i|a_i)ds_i}_{\text{direct effect}}$$

$$+ \underbrace{\int_{-\infty}^{\infty} \left(\frac{\partial}{\partial \underline{s}}\mathbb{P}\left(a_i + e_i + \tilde{\epsilon}_i^q \geq t_k | a_i, s_i\right)\frac{\partial \underline{s}(t_k)}{\partial t_k} + \frac{\partial}{\partial \overline{s}}\mathbb{P}\left(a_i + e_i + \tilde{\epsilon}_i^q \geq t_k | a_i, s_i\right)\frac{\partial \overline{s}(t_k)}{\partial t_k}\right)\phi(s_i|a_i)ds_i}_{\text{incentive effect}}, \quad (6)$$

where we have $\frac{\partial \underline{s}(t_k)}{\partial t_k} = \frac{\partial \overline{s}(t_k)}{\partial t_k} = \frac{\sigma_k^2 + \sigma_s^2}{\sigma_k^2}$ from Lemma 1.

This incentive effect can be either positive or negative. Proposition 1 studies its sign, and also shows how the level of the threshold that maximizes this incentive effect depends on the individual's talent.

Proposition 1 *For an individual with talent a_i, the incentive effect is positive if and only if t_k is sufficiently low. Moreover, the level of the threshold t_k that maximizes the incentive effect for an individual with talent a_i is increasing in a_i.*

As established in Lemma 1, an increase in the threshold t_k increases the effort bounds $\underline{s}$ and $\overline{s}$. This increases the probability that an individual with talent a_i exerts effort if and only if $\phi(s_i|a_i)$, the conditional density of s_i, is larger at $s_i = \overline{s}$ than at $s_i = \underline{s}$. In turn, this is the case if and only if the threshold is sufficiently low. In this case, the "stimulant effect" of a higher threshold (effort inducement for individuals whose signal is close to $s_i = \overline{s}$) dominates the "discouragement effect" (effort withdrawal for individuals whose signal is close to $s_i = \underline{s}$). The crucial assumption underlying this result is that the talent distribution is single-peaked. In sum, higher thresholds tend to elicit more effort from more talented individuals, but less effort from less talented individuals.

5 The talent pool

We now consider the measure of appointed individuals from a given group. In particular, we study the effect of the threshold on the pipeline of eligible CEO candidates. The firm CEO is selected among appointed individuals whose performance is higher than a target level of ability denoted by $\bar{a}$. With a continuum $[0, 1]$ of candidates, the ex-ante probability for a member of group k to be appointed is also the measure of appointed individuals from this group.

5.1 Perfect information

In this subsection, we focus on the case when individual abilities are common knowledge, i.e. $\sigma_s = 0$ and $\sigma_q = 0$. By abstracting from uncertainty, these results also serve as stepping stones to better understand the effects at play in more complex settings.

Proposition 2 studies analytically the effect of the threshold on the measure of appointed candidates whose ability is sufficiently high as to be CEO-eligible.

Proposition 2 *When individuals know their talent and screening is perfect, the measure of appointed individuals with ability larger than $\bar{a}$ is increasing in the threshold t_k if $t_k < \bar{a}$.*

When individual abilities are known and the targeted level of ability $\bar{a}$ is above the appointment threshold, there is no direct effect. Indeed, with a perfect appointment process, an individual exceeds the threshold if and only if her ability is sufficiently high, so that raising the threshold does not result in the unintended rejection of high-ability candidates. In addition, the incentive effect is then always nonnegative because we focus on individuals whose ability is above the threshold, and the appointment process is perfect. In this case, the stimulant effect is relevant, whereas the discouragement effect is not. In sum, in the case considered in Proposition 2, raising the appointment threshold always increases the measure of appointed individuals with sufficiently high ability.

This analytical result implies that, when there is no uncertainty about individual ability, a "pipeline" of very high ability individuals from a particular group is best cultivated by setting a very high appointment threshold. Indeed, this policy maximizes incentives for individuals with high talent to exert effort, which thickens the right tail of the ability distribution.

This result relies on strong assumptions. When individuals are imperfectly aware of their own talent ($\sigma_s > 0$), setting a very high threshold could backfire because some individuals with high talent will not exert effort. Likewise, when the appointment process is not perfect ($\sigma_q > 0$), setting a very high threshold could backfire because some candidates with high ability will not be appointed. We explore these possibilities in the next subsection.

5.2 Imperfect information

We now study the case when individuals imperfectly know their own talent ($\sigma_s > 0$) and the screening process is imperfect ($\sigma_q > 0$). Allowing for $\sigma_s > 0$ and $\sigma_q > 0$ introduces two additional effects.

First, the less precise the signal that individuals receive about their talent (high σ_s), the more they rely on their prior belief when making their effort decision. In particular, individuals whose talent is significantly higher than the mean $\mu_k = 0$ in their group should in principle exert effort when the recruitment threshold is high. However, if they are not aware of their high talent, they will tend to make effort decisions similar to the average individual in their group, and thus have the strongest propensity to exert effort when the threshold is intermediate rather than high. In this case, high thresholds will discourage effort even among talented individuals.

Second, a more noisy appointment process results in more mistakes, including appointments of low ability individuals and rejections of high ability individuals. When the appointment process is substantially noisy (high σ_q), appointment decisions are largely random, which discourages effort.[10] Then, maximizing the measure of appointed individuals with ability higher than a threshold is achieved with a low threshold. On the contrary, when the appointment process is high-quality (low σ_q), increasing the threshold allows the organization to reject more individuals with low ability (direct effect), and it also encourages high ability individuals to exert effort (incentive effect) – all the more that these individuals receive a precise signal about their type.

We now consider the measure of individuals from group k who are appointed to entry-level positions *and* whose ability is higher than a cutoff $\bar{a}$. This measure is given by integrating over a_i such that $a_i + e_i \geq \bar{a}$ the measure of appointed individuals with talent a_i derived in equation (5):

$$
\begin{aligned}
\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right) \quad = \quad & \int_{\bar{a}}^{\infty} \int_{-\infty}^{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\right) \phi(s_i|a_i)\varphi(a_i) ds_i da_i \\
& + \int_{\bar{a}-\bar{e}}^{\infty} \int_{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i - \bar{e}\right)\right)\right) \phi(s_i|a_i)\varphi(a_i) ds_i da_i \\
& + \int_{\bar{a}}^{\infty} \int_{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\infty} \left(1 - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\right) \phi(s_i|a_i)\varphi(a_i) ds_i da_i \qquad (7)
\end{aligned}
$$

where φ is the PDF of the standard normal distribution. The integrals on the first and third lines account for individuals who do not exert effort, whose ability is simply equal to their talent. The integrals on the second line account for individuals who exert effort, whose ability is equal to the sum of their talent and effort. From Lemma 1, we know that only individuals who receive a signal s_i in-between the "effort bounds" will exert effort.

We normalized the mean of the ability distribution to zero and its variance to one. We now consider the calibration of other parameters, which will matter for numerical results.

First, the cost of effort must be strictly less than the benefit B of being appointed, otherwise no one would exert effort. Moreover, because of imperfect information, the cost of effort cannot be too high. For some parameter values considered, the highest cost of effort such that a positive measure of individuals exerts effort is $c = 0.31$. A cost of effort which is too low would simply result in the vast majority of individuals exerting effort, which may not be very theoretically interesting or practically relevant. Accordingly, we assume a cost of effort of $c = 0.3$ throughout. We also set the effect of effort on ability equal to one standard deviation of the ability distribution: $\bar{e} = \sigma_k = 1$. This similar magnitude implies that both talent and effort are important to achieve a very high level of ability.

[10]Considering the effort incentive constraint for a given s_i in equation (8) in the Appendix, it is not satisfied if σ_q is sufficiently large. This means that an individual optimally decides not to exert effort if appointments are sufficiently random.

Second, meta-analyses find that the best combination of test scores used by firms for screening explains between 26% to 42% of the variance of actual work performance, a fraction which has been increasing over time (see Schmidt, Oh, and Shaffer (2016), which updates the seminal paper of Schmidt and Hunter (1998)). This suggests that σ_q should be $\sigma_q \approx 0.8$.[11]

Third, meta-analyses of the correlation ρ between students' self-assessments of their own ability and the marks they actually receive suggest that $\rho \approx 0.45$ (Falchikov and Boud (1989), León, Panadero, and García-Martínez (2023)),[12] which given $\sigma_q \approx 0.8$ suggests that $\sigma_s \approx 1.42$. When individuals are more educated or receive feedback, the correlation coefficient is typically around $\rho \approx 0.5$ or higher (Ackerman, Kanfer, and Goff (1995), Ober et al. (2022), Atrash, Katz-Leurer, and Shahar (2023), Patzl, Oberleiter, and Pietschnig (2024)), which given $\sigma_q \approx 0.8$ suggests that $\sigma_s \approx 1.20$.[13] Accordingly, we set $\sigma_s = 1.20$ in our baseline numerical computations, and we repeat the same computations with different values for σ_s in the section below and in Appendix B.

We consider the effect of the threshold on the measure of appointed individuals with high ability. Figure 1 depicts the measure of individuals from group k who get appointed ($q_i \geq t_k$) and whose ability is higher than a cutoff $\bar{a}$, as a function of the threshold t_k. The target $\bar{a}$, the minimum level of ability required to be selected as CEO, can be viewed as a measure of the meritocracy of the CEO selection process: the higher $\bar{a}$, the more ability matters; the lower it is, the more other factors matter.[14] Figure 1 presents numerical computations based on our baseline parameter values, as well as additional computations that either simply assume $\sigma_q = 1$ and $\sigma_s = 1$, or assume very precise information about individual abilities (see Appendix B for additional results based on alternative parameter values). We can draw two main lessons from these numerical computations.

First, increasing the appointment threshold for a group can substantially increase the measure of appointed individuals with high ability: the function $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ is increasing in the appointment threshold t_k when the threshold is low.[15] Thus, the forces at play in Proposition 2 remain relevant and can be economically significant even with $\sigma_s > 0$ and $\sigma_q > 0$.

Second, these numerical computations also show that increasing the threshold reduces the

[11] Back of the envelope calculations that abstract from effort show that: $R^2 = 1 - \frac{var(\tilde{\epsilon}_i^q)}{var(\tilde{a}_i)} = 1 - \sigma_q^2$. With $R^2 = 0.42$, this gives $\sigma_q = 0.76$; with $R^2 = 0.26$, this gives $\sigma_q = 0.86$.

[12] Falchikov and Boud (1989) find a distribution of correlation coefficients across studies that peaks at $\rho = 0.45$, while León, Panadero, and García-Martínez (2023) find a Fisher z-transformation of 0.472, corresponding to a correlation coefficient of 0.44.

[13] Using the updated belief about own ability in equation (2), back of the envelope calculations that abstract from effort show that a correlation of ρ means that: $\rho = \text{corr}\left(\mathbb{E}[\tilde{a}_i|\tilde{s}_i], q_i\right) = \text{corr}\left(\frac{\sigma_k^2}{\sigma_k^2+\sigma_s^2}(\tilde{a}_i + \tilde{\epsilon}_i^s), \tilde{a}_i + \tilde{\epsilon}_i^q\right) = \frac{1}{\sqrt{1+\sigma_s^2}\sqrt{1+\sigma_q^2}}$.

[14] These other factors may include connections both inside and outside the firm, including linkages to the previous CEO and linkages to directors (Wiersema, Nishimura, and Suzuki (2018), Bandiera et al. (2024)).

[15] This is only true when $\overline{a}$ is sufficiently high. For example, for $\overline{a} \to -\infty$, this function would be strictly decreasing in the appointment threshold at any point.

measure of appointed individuals with high ability when the threshold is high: the measure $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ is decreasing in the appointment threshold t_k when it is sufficiently high. This can happen even with $t_k < \bar{a}$, which is in contrast to the result in Proposition 2. When there is uncertainty about individual abilities, the appointment threshold that maximizes the measure of high-ability employees is not necessarily as high as the targeted level of ability.

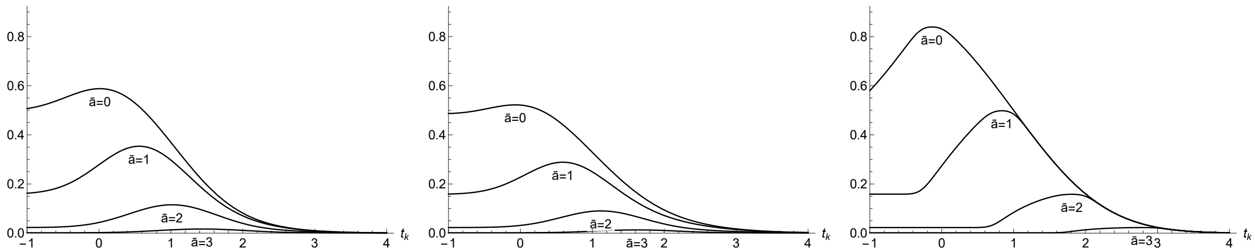

Figure 1: Measure of appointed high-ability employees $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ as a function of the appointment threshold t_k for $\mu_k = 0$, $\sigma_k = 1$, $\overline{e} = 1$, $c = 0.3$.
Left panel: $\sigma_s = 1.2$, $\sigma_q = 0.8$; middle panel: $\sigma_s = 1$, $\sigma_q = 1$; right panel: $\sigma_s = 0.1$, $\sigma_q = 0.1$.

Next, we consider the threshold that maximizes the measure of appointed individuals with ability above the minimum level necessary to be selected as CEO. Figure 2 depicts the threshold t_k^* that maximizes the measure of appointed individuals with ability above $\bar{a}$, for a range of values of $\bar{a}$, as computed numerically. In all cases, this threshold is increasing in the minimum level of ability $\bar{a}$. Intuitively, changing the threshold shifts the subset of individuals who exert effort. When the threshold increases, less talented individuals no longer exert effort (the discouragement effect), whereas more talented individuals start exerting effort (the stimulant effect). Since the talent distribution is single-peaked with an upper tail that gets thinner and thinner, increasing the threshold from a high level involves a tradeoff between eliciting effort from more talented individuals and eliciting effort from fewer individuals. As the cutoff $\bar{a}$ increases, is becomes even more important to elicit effort from more talented individuals, so that the threshold that maximizes the measure increases.

It is noticeable that the sensitivity of the threshold t_k^* to the minimum level of ability $\bar{a}$ is increasing in the precision $\frac{1}{\sigma_s^2}$ of individuals' signals about their own talent. Indeed, when these signals are more precise, higher thresholds are more effective at eliciting effort from more talented individuals. This effect can also be seen in Appendix A, which displays the effect of the threshold on the ability distribution of appointed individuals when information about individual abilities is very precise and when it is not.

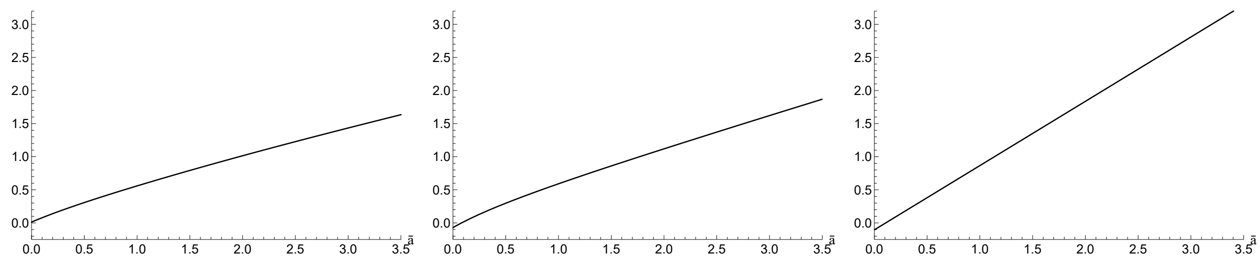

Figure 2: Appointment threshold t_k^* that maximizes the measure of appointed high-ability employees $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ as a function of $\bar{a}$ for $\mu_k = 0$, $\sigma_k = 1$, $\overline{e} = 1$, $c = 0.3$. Left panel: $\sigma_s = 1.2$, $\sigma_q = 0.8$; middle panel: $\sigma_s = 1$, $\sigma_q = 1$; right panel: $\sigma_s = 0.1$, $\sigma_q = 0.1$.

The effect of the appointment threshold on the probability for a member of a group to become CEO is relative: it depends not only on the threshold applied to this group, but also on the threshold applied to other groups. In Figure 3, we consider two groups, A and B, subject to different appointment thresholds, t_A and t_B, and we ask whether the firm CEO is more likely to be from group A or from group B, depending on the thresholds. As long as the appointment thresholds are sufficiently low, we find that the CEO is more likely to be from the group with the highest appointment threshold. On the contrary, when the thresholds are sufficiently high, the CEO is more likely to be from the group with the lowest appointment threshold. Moreover, the more meritocratic the firm (i.e. the higher $\bar{a}$), the larger is the subset of thresholds such that the CEO is more likely to be from the group with the highest threshold. For example, for $t_A = 1.25$ and $t_B = 1.5$, the CEO is more likely to be from group A if $\bar{a} = 2$, but she is more likely to be from group B is $\bar{a} = 3$.

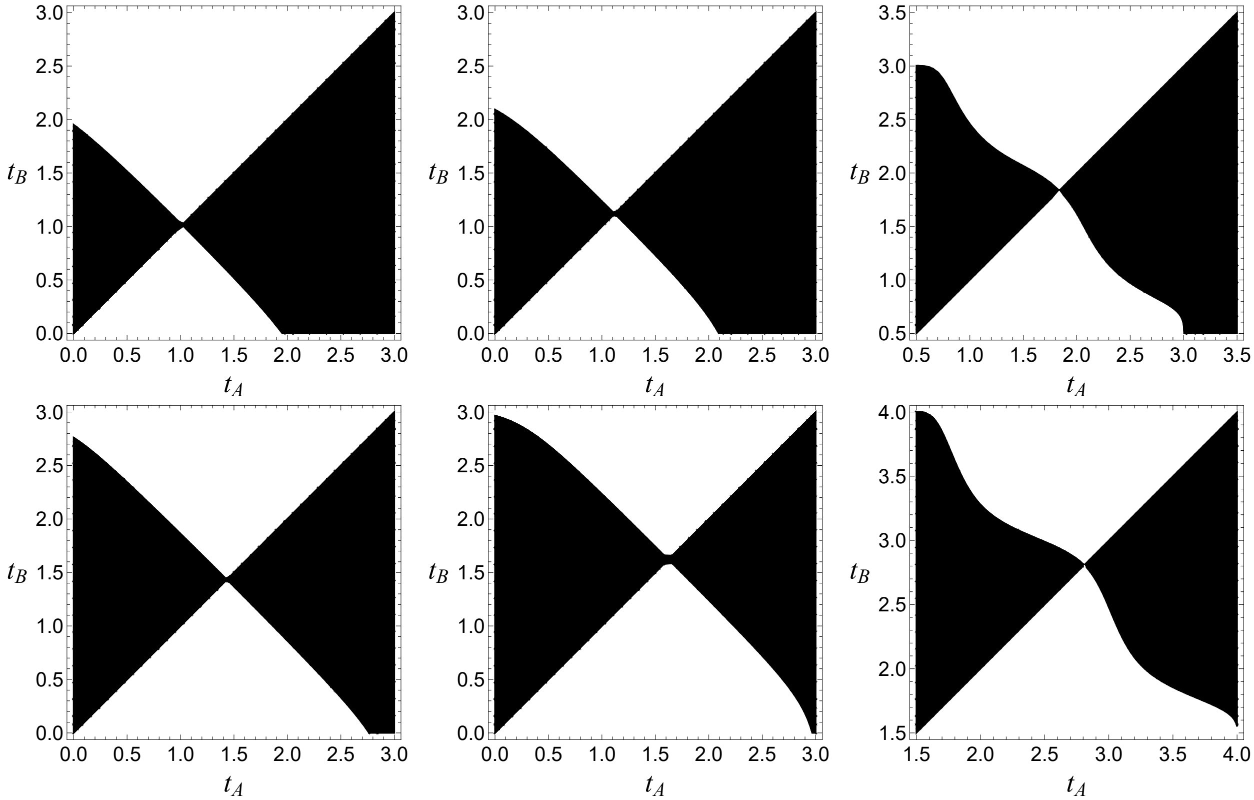

Figure 3: The black area is the subset of group-contingent appointment thresholds $\{t_A, t_B\}$ such that the measure of appointed high-ability employees $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ is higher for group B than for group A, for $\mu_k = 0$, $\sigma_k = 1$, $\overline{e} = 1$, $c = 0.3$.
First row: $\bar{a} = 2$. Second row: $\bar{a} = 3$. On each row: left panel: $\sigma_s = 1.2$, $\sigma_q = 0.8$; middle panel: $\sigma_s = 1$, $\sigma_q = 1$; right panel: $\sigma_s = 0.1$, $\sigma_q = 0.1$.

Finally, to quantify the magnitude of the effects at play, we present numerical examples that involve two firms with the same group-contingent appointment thresholds (i.e. biases are systematic) but different levels of meritocracy. We consider two different cases. Example 1 relies on our baseline parameter values, and considers moderate appointment thresholds. Example 2 considers the hypothetical case when abilities are very precisely estimated and appointment thresholds are high.

Example 1 *Individual talent is normally distributed with mean $\mu_k = 0$ and variance $\sigma_k^2 = 1$. The cost of effort is $c = 0.3$, and the effect of effort on ability is equal to one standard deviation of the talent distribution: $\overline{e} = 1$. There are two firms. The first is highly meritocratic: it only consider very high-ability individuals as potential CEOs, with $\overline{a} = 3$. This is two standard deviations above the mean of the ability distribution for individuals who exert effort. By contrast, the second firm considers a broader pool of individuals as potential CEOs, with $\overline{a} = 2$. This is one standard deviation above the mean of the ability distribution for individuals who exert* *effort. There are two groups in the population, group A and group B, with the same measure of individuals from each group. The appointment threshold is biased in favor of individuals from group A: $t_A = 1.0$ and $t_B = 1.5$: the difference in thresholds across the two groups is half a standard deviation of the talent distribution. Using our baseline calibration for σ_s and σ_q, we have:*

$\sigma_s = 1.2$, $\sigma_q = 0.8$	*Proportion of individuals from group B in the firm*	*Probability that the CEO is from group B*
More meritocratic firm: $\bar{a} = 3$	*34.4%*	*55.7%*
Less meritocratic firm: $\bar{a} = 2$	*34.4%*	*43.2%*

Example 2 *Consider the same parameter values as in Example 1 except that $\sigma_s = 0.1$, $\sigma_q = 0.1$, $t_B = 2.5$, and $t_A = 2.0$. Then we have:*

$\sigma_s = 0.1$, $\sigma_q = 0.1$	*Proportion of individuals from group B in the firm*	*Probability that the CEO is from group B*
More meritocratic firm: $\bar{a} = 3$	*29.7%*	*65.7%*
Less meritocratic firm: $\bar{a} = 2$	*29.7%*	*30.0%*

In cases considered in Examples 1 and 2, we find that the probability that the CEO is from the group with the higher appointment threshold is higher than this group representation in the firm. Intuitively, its talented members will have stronger effort incentives than talented members from the other group, but its less talented members will most likely not get appointed.

However, this does not imply that appointment biases against a group will automatically increase the probability that the firm CEO is a member of this group. This can be for two reasons. First, in the less meritocratic firm, the probability that the CEO is from group B is less than one-half in both Examples, even though group B accounts for one-half of the population. These firms do not require their CEO to have very high ability, so that the probability that the CEO is from a particular group mostly depends on the measure of appointed individuals from this group. Second, if individuals receive a very noisy signal about their own talent, talented individuals from a group subject to a high appointment threshold will fail to exert effort because they put a lot of weight on their prior belief μ_k about their talent. Formally:

Remark 1 *If t_k and σ_s are sufficiently high, there is no signal realization s_i such that an individual from group k will optimally decide to exert effort.*

When no one exerts effort, the "incentive effect" of changing the threshold (as described in section 4.2) does not operate, so that a higher threshold simply reduces the number of individuals from a group recruited by the firm. For example, using our baseline parameter values ($\mu_k = 0$, $\sigma_k = 1$, $\bar{e} = 1$, $c = 0.3$, $\sigma_q = 0.8$), no individual from group k exerts effort if $\sigma_s > 7.61$.

6 Discussion and implications

The model developed in this paper highlights an important tradeoff in recruitment practices. Algorithmic or procedural biases that result in higher appointment thresholds for some demographic groups can have opposing effects at different levels of the corporate hierarchy. While a higher threshold reduces the representation of that group at the entry level, it can increase the share of high-ability individuals from the same group at the top, especially in meritocratic firms where promotions are tightly linked to individual ability.

This finding has direct implications for the use of AI in recruitment. Many hiring algorithms are trained on historical data that embed prior human biases and thus systematically vary thresholds across groups. Such biases are often viewed as detrimental, prompting firms to put in place bias-correction mechanisms. However, the model suggests that eliminating algorithmic bias without considering incentive effects may have unintended long-term consequences: it can weaken incentives for talented individuals in underrepresented groups to invest in human capital, potentially reducing top-management diversity in meritocratic organizations.

More broadly, the results emphasize that diversity initiatives and AI auditing practices cannot be evaluated solely based on short-term hiring outcomes. Algorithmic systems influence human capital investment long before promotion decisions occur. More demanding selection processes can strengthen motivation and effort among talented individuals, especially in meritocratic organizations. However, this mechanism depends critically on firm-level meritocracy and the quality of information available. In less meritocratic firms or when individuals have poor feedback about their own talent, higher thresholds suppress participation at all levels.

As AI-driven recruitment becomes more prevalent, improvements in screening accuracy and in individuals' ability to assess their own talent will effectively reduce the values of the parameters σ_q and σ_s, which are the key frictions in the model. This technological progress may amplify the tradeoff between diversity at junior and senior levels.

The results also imply that the structure of hiring and promotion incentives can have substantial effects on the distribution of talent within the firm and, consequently, on firm value. In well-governed firms, even small improvements in the quality of the internal talent pool can translate into substantial value creation, especially in large firms where CEO ability matters more (Gabaix and Landier (2008)).

Finally, the model generates testable predictions. Most notably, for groups that face discrimination at the hiring stage, their representation at the CEO level should be positively correlated with measures of meritocracy across firms – as proxied by governance quality or shareholder orientation. To take into account firm fixed effects in group representation, an empirical test could alternatively consider senior level representation relative to junior level representation across firms.

7 Conclusion

In this paper, we study the implications of group-dependent appointment thresholds for human capital accumulation decisions optimally made by individuals, for entry-level appointments, and for CEO selection. A higher appointment threshold for a group motivates more talented individuals in this group to exert effort. As a result, even though it decreases entry-level representation, it can increase the number of high-ability individuals from this group who are appointed. In turn, assuming that eligibility for the CEO position requires an ability level exceeding a cutoff, a higher appointment threshold for a particular group may increase the probability that the firm's CEO is selected from that group. We find that this tradeoff arises when individuals are well-informed about their own talent and when firms are highly meritocratic, in the sense that the aforementioned cutoff is high. In summary, systematic biases that impede the appointment of individuals from a group can thicken the upper tail of the ability distribution in this group, which enlarges the pool of candidates from which top management is selected.

References

Ackerman, P.L., Kanfer, R. and Goff, M., 1995. Cognitive and noncognitive determinants and consequences of complex skill acquisition. *Journal of Experimental Psychology: Applied*, 1(4), 270.

Adams, R.B., Ferreira, D., 2009. Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics*, 94(2), 291-309.

Atrash, H., Katz-Leurer, M. and Shahar, G., 2023. The effect of self-assessment on student competence in physiotherapy clinical training: a randomized controlled trial. *BMC Medical Education*, 23(1), 780.

Bandiera, O., Kotia, A., Lindenlaub, I., Moser, C. and Prat, A., 2024. Meritocracy across countries. NBER working paper w32375.

Bertrand, M. and Duflo, E., 2017. Field experiments on discrimination. *Handbook of Economic Field Experiments*, 1, 309-393.

Blair, P.Q., Chung, B.W., 2021. A model of occupational licensing and statistical discrimination. *American Economic Association Papers and Proceedings*, 111, 201-205.

Blau, F.D., Kahn, L.M., 2017. The gender wage gap: extent, trends, and explanations. *Journal of Economic Literature* 55(3), 789-865.

Burkart, M., Panunzi, F. and Shleifer, A., 2003. Family firms. *Journal of Finance*, 58(5), 2167-2201.

Cannady, M.A., Greenwald, E., Harris, K.N., 2014. Problematizing the STEM Pipeline Metaphor: Is the STEM Pipeline Metaphor Serving Our Students and the STEM Workforce? *Science Education*, 98, 443-460.

Castleman, T., Allen, M., 1998. The 'Pipeline Fallacy' and Gender Inequality in Higher Education Employment. *Policy, Organisation and Society*, 15(1), 23-44.

Chen, Z., 2023. Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications*, 10(1), 1-12.

Chowdhury, S.M., Esteve-González, P. and Mukherjee, A., 2023. Heterogeneity, leveling the playing field, and affirmative action in contests. *Southern Economic Journal*, 89(3), 924-974.

Coate, S., Loury, G.C., 1993. Will Affirmative-Action Policies Eliminate Negative Stereotypes? *American Economic Review*, 83 (5), 1220-1240.

Cotton, C., Hickman, B.R., List, J.A., Price, J., Roy, S., 2025. Why Don't Struggling Kids Do Their Homework? Disentangling Motivation and Study Productivity in a Model of Human Capital Investment. *Journal of Political Economy*, forthcoming.

Cook, A., Glass, C., 2014. Above the glass ceiling: When are women and racial/ethnic minorities promoted to CEO? *Strategic Management Journal*, 35, 1080-1089.

Cowgill, B. and Tucker, C.E., 2020. Algorithmic fairness and economics. Working paper, Columbia.

Cziraki, P., Jenter, D., 2022. The Market for CEOs. Working paper, University of Toronto.

Dezsö, C.L., Ross, D.G., Uribe, J., 2016. Is there an implicit quota on women in top management? A large-sample statistical analysis. *Strategic Management Journal*, 37, 98-115.

Espenshade, T.J., & Radford, A.W., 2009. *No longer separate, not yet equal: Race and class in elite college admission and campus life*. Princeton, NJ: Princeton University Press.

Fabris, A., Baranowska, N., Dennis, M.J., Graus, D., Hacker, P., Saldivar, J., Zuiderveen Borgesius, F. and Biega, A.J., 2025. Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Transactions on Intelligent Systems and Technology*, 16(1), 1-54.

Falchikov, N. and Boud, D., 1989. Student self-assessment in higher education: A meta-analysis. *Review of Educational Research*, 59(4), 395-430.

Fang, H., Moro, A., 2011. Chapter 5 – Theories of Statistical Discrimination and Affirmative Action: A Survey. *Handbook of Social Economics*, 1, North-Holland, 133-200.

Fernandez-Mateo, I., Fernandez, R.M., 2016. Bending the Pipeline? Executive Search and Gender Inequality in Hiring for Top Management Jobs. *Management Science*, 62(12), 3636-3655.

Ferreira, D., Nikolowa, R. and Pikulina, E., 2025. Biased Promotions. Working paper, London School of Economics.

Fershtman, D., Pavan, A., 2021. Soft Affirmative Action and Minority Recruitment. *American Economic Review: Insights*, 3(1), 1-18.

Fryer, R.G., 2007. Belief flipping in a dynamic model of statistical discrimination. *Journal of Public Economics*, 91, 1151-1166.

Fryer Jr., R.G., Loury, G.C., 2013. Valuing Diversity. *Journal of Political Economy*, 121(4), 747-774.

Goldin, C. and Rouse, C., 2000. Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, 90(4), 715-741.

He, L., Whited, T., 2023. Underrepresentation of Women CEOs. Working paper, Erasmus University.

Helfat, C.E., Harris, D., Wolfson, P.J., 2006. The Pipeline to the Top. *Academy of Management Perspectives*, 20, 42-64.

Huesmann, K., Ndiaye, Y.S., Waibel, C., Wiesen, D., 2025. How the Design of Ranking Systems and Ability Affects Physician Effort. *Management Science*, forthcoming.

Kaplan, S.N. and Minton, B.A., 2012. How has CEO turnover changed? *International review of Finance*, 12(1), 57-87.

Köchling, A. and Wehner, M.C., 2020. Discriminated by an algorithm: a systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13(3), 795-848.

Krishna, K., Tarasov, A., 2016. Affirmative Action: One Size Does Not Fit All. *American Economic Journal: Microeconomics*, 8(2), 215-252.

Larcker, D.F., Tayan, B., 2020. Diversity in the C-Suite. The Dismal State of Diversity Among Fortune 100 Senior Executives. *Rock Center for Corporate Governance*, Stanford, CGRP82.

León, S.P., Panadero, E. and García-Martínez, I., 2023. How accurate are our students? A meta-analytic systematic review on self-assessment scoring accuracy. *Educational Psychology Review*, 35(4), 106.

Matsa, D.A., Miller, A.R., 2011. Chipping Away at the Glass Ceiling: Gender Spillovers in Corporate Leadership. *American Economic Review*, 101(3), 635-639.

Moro, A., Norman, P., 2004. A general equilibrium model of statistical discrimination. *Journal of Economic Theory*, 114(1), 1-30.

Moro, A., Norman, P., 2003. Affirmative action in a competitive economy. *Journal of Public Economics*, 87, 567-594.

Neely B.H., Lovelace J.B., Cowen A.P., Hiller N.J., 2020. Metacritiques of Upper Echelons Theory: Verdicts and Recommendations for Future Research. *Journal of Management*, 46(6), 1029-1062.

Nielsen, S., 2010. Top Management Team Diversity: A Review of Theories and Methodologies. *International Journal of Management Reviews*, 12, 301-316.

Ober, T.M., Hong, M.R., Carter, M.F., Brodersen, A.S., Rebouças-Ju, D., Liu, C. and Cheng, Y., 2022. Are high school students accurate in predicting their AP exam scores?: Examining inaccuracy and overconfidence of students' predictions. *Assessment in Education: Principles, Policy & Practice*, 29(1), 27-50.

Pagano, M. and Picariello, L., 2025. Corporate governance, meritocracy, and careers. *Review of Finance*, 29(2), 349-379.

Patzl, S., Oberleiter, S. and Pietschnig, J., 2024. Mirror, Mirror on the Wall: A Meta-Analysis on the Validity of Self-Assessed Intelligence through the Lens of the Multiverse. *Journal of Intelligence*, 12(9), 81.

Peck, J.R., 2017. Can Hiring Quotas Work? The Effect of the Nitaqat Program on the Saudi Private Sector. *American Economic Journal: Economic Policy*, 9(2), 316-47.

Pikulina, E.S. and Ferreira, D., 2025. Subtle discrimination. *Journal of Finance*, forthcoming.

Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K., 2020, January. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 469-481.

Schaede, U. and Mankki, V., 2022. Quota vs Quality? Long-term gains from an unusual gender quota. Working paper, University of Zurich.

Schmidt, F.L. and Hunter, J.E., 1998. The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological Bulletin*, 124(2), 262.

Schmidt, F.L., Oh, I.S. and Shaffer, J.A., 2016. The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 100 years. Working paper, University of Iowa.

Sethi, R., Somanathan, R., 2018. Meritocracy in the Face of Group Inequality. Working Paper, Columbia University.

Sethi, R. and Somanathan, R., 2023. Meritocracy and representation. *Journal of Economic Literature*, 61(3), 941-957.

Smith N., Smith, V., Verner, M., 2013. Why are So Few Females Promoted into CEO and Vice President Positions? Danish Empirical Evidence, 1997-2007. *ILR Review*, 66(2), 380-408.

Spence, M., 1973. Job market signaling. *Quarterly Journal of Economics*, 87(3), 355-374.

Wiersema, M.F., Nishimura, Y. and Suzuki, K., 2018. Executive succession: The importance of social capital in CEO appointments. *Strategic Management Journal*, 39(5), 1473-1495.

Appendix

A Ability distribution

Figure 4 depicts the distribution of the ability of appointed individuals for four different levels of the appointment threshold t_k. These graphs allow to better understand how different values of t_k affect effort decisions and the ability distribution of appointed individuals. First, a low threshold ($t_k = 0$) results in a very high measure of appointed individuals with average and moderately high ability, but a very low measure of appointed individuals with very high ability (above $\mu_k + 3\sigma_k = 3$). Second, a high threshold ($t_k \geq 3$) results in a low measure of appointed individuals, but a relatively high measure of appointed individuals with very high ability when σ_s and σ_q are low (see the bottom row in Figure 4).

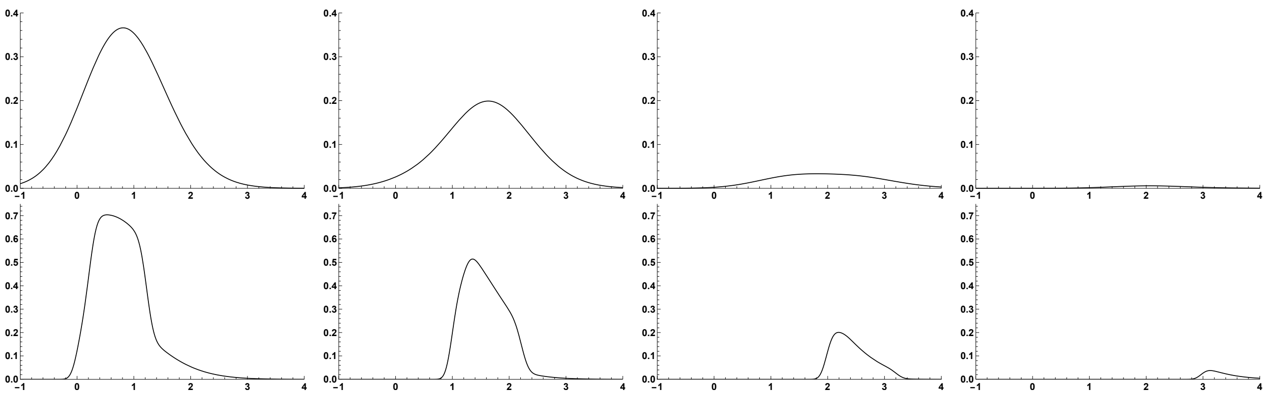

Figure 4: Measure of the ability $\tilde{a}_i + e_i$ of appointed individuals, for $\mu_k = 0$, $\sigma_k^2 = 1$, $\overline{e} = 1$, $c = 0.3$.
Top row: $\sigma_q = 0.8$, $\sigma_s = 1.2$, and from left to right: $t_k = 0$, $t_k = 1$, $t_k = 2$, $t_k = 3$.
Bottom row: $\sigma_q = 0.1$, $\sigma_s = 0.1$, and from left to right: $t_k = 0$, $t_k = 1$, $t_k = 2$, $t_k = 3$.

B Alternative parameter values

This section presents numerical findings based on alternative parameter values for robustness purposes. The main text already presented findings with lower values for σ_q and σ_s. Figure 5 presents findings with higher values for these two key parameters.

In the case $\sigma_q = 1.5$, $\sigma_s = 1$, there does not exist any signal realization such that an individual exerts effort when $c = 0.3$ (the benefits are always too low relative to the costs). Accordingly, in this case we set a lower cost of effort: $c = 0.25$.

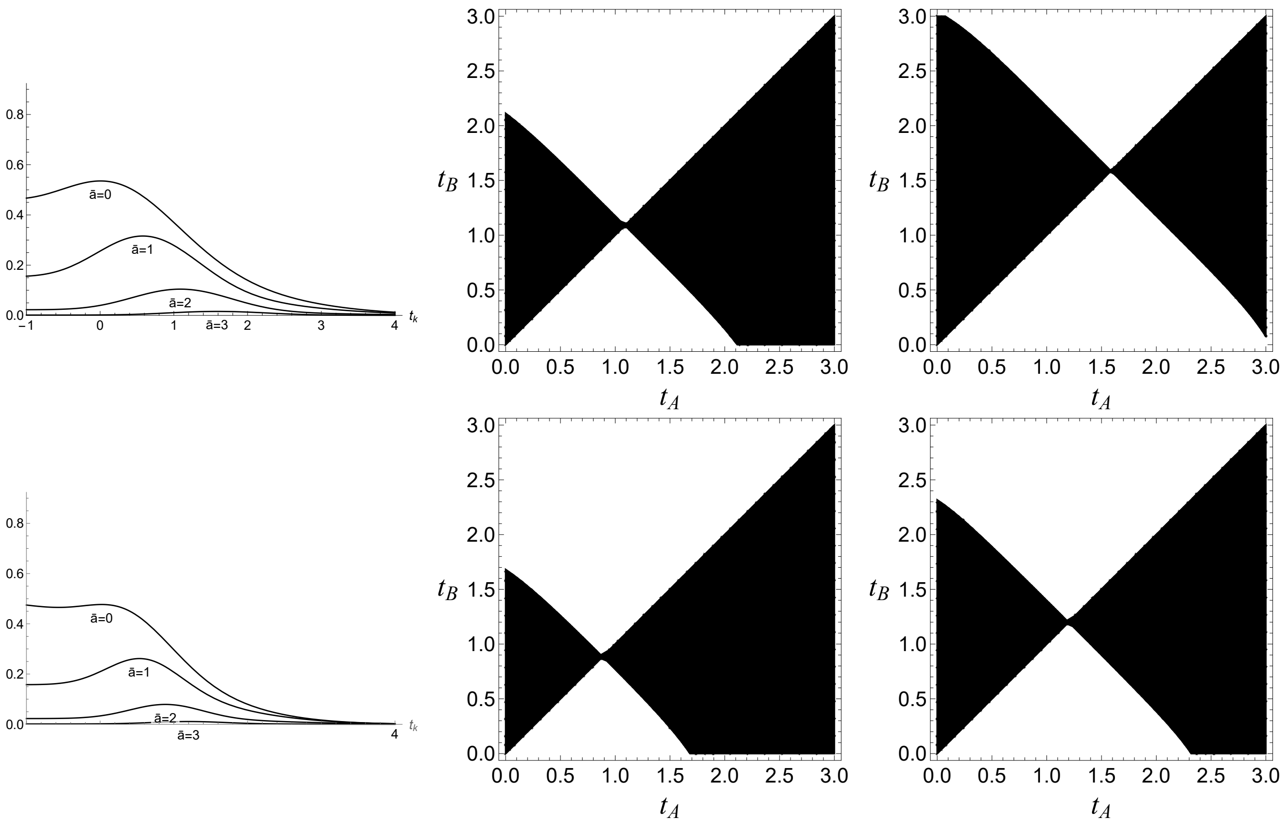

Figure 5: Left panel: $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ as a function of t_k for different values of $\bar{a}$. Middle and right panels: the black area is the subset of $\{t_A, t_B\}$ such that the measure $\mathbb{P}\left(\tilde{q}_i \geq t_k \cap \tilde{a}_i + e_i \geq \bar{a}\right)$ is higher for group B than for group A; middle panel: $\bar{a} = 2$; right panel: $\bar{a} = 3$. In all panels, $\mu_k = 0$, $\sigma_k^2 = 1$, $\overline{e} = 1$. Top row: $\sigma_q = 1.5$, $\sigma_s = 1$, $c = 0.25$. Bottom row: $\sigma_q = 1$, $\sigma_s = 1.5$, $c = 0.3$.

C Proofs

Proof of Lemma 1:

Conditional on the signal s_i, the random variable $\tilde{a}_i + \tilde{\epsilon}_i^q$ is normally distributed with mean

$$\mu_k + \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]$$

and variance

$$\Sigma^2 \equiv \frac{\sigma_k^2 \sigma_s^2}{\sigma_k^2 + \sigma_s^2} + \sigma_q^2.$$

Based on her belief about her own talent described in equations (2) and (3), individual i optimally exerts effort if and only if:

$$\begin{aligned}
& \mathbb{P}(\tilde{q}_i \geq t_k | e_i = \overline{e}, s_i) - \mathbb{P}(\tilde{q}_i \geq t_k | e_i = 0, s_i) \geq c \\
\Leftrightarrow \quad & \mathbb{P}(\tilde{a}_i + \tilde{\epsilon}_i^q \geq t_k - \overline{e} | s_i) - \mathbb{P}(\tilde{a}_i + \tilde{\epsilon}_i^q \geq t_k | s_i) \geq c \\
\Leftrightarrow \quad & \mathbb{P}\left(\tilde{\varepsilon} \geq \left(\sqrt{\frac{\sigma_k^2 \sigma_s^2}{\sigma_k^2 + \sigma_s^2}} + \sigma_q^2\right)^{-1} \left(t_k - \overline{e} - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) \\
& -\mathbb{P}\left(\tilde{\varepsilon} \geq \left(\sqrt{\frac{\sigma_k^2 \sigma_s^2}{\sigma_k^2 + \sigma_s^2}} + \sigma_q^2\right)^{-1} \left(t_k - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) \geq c \\
\Leftrightarrow \quad & 1 - \Phi\left(\left(\sqrt{\frac{\sigma_k^2 \sigma_s^2}{\sigma_k^2 + \sigma_s^2}} + \sigma_q^2\right)^{-1} \left(t_k - \overline{e} - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) \\
& -1 + \Phi\left(\left(\sqrt{\frac{\sigma_k^2 \sigma_s^2}{\sigma_k^2 + \sigma_s^2}} + \sigma_q^2\right)^{-1} \left(t_k - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) \geq c, \qquad (8)
\end{aligned}$$

where the random variable $\tilde{\varepsilon}$ follows a standard normal distribution with CDF Φ. The LHS of equation (8), which measures the expected benefits from effort, is positive since $\Phi' > 0$. Let s^* be implicitly defined by:

$$\Phi\left(\Sigma^{-1}\left(t_k - \mu_k - \tfrac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s^* - \mu_k\right]\right)\right) - \Phi\left(\Sigma^{-1}\left(t_k - \overline{e} - \mu_k - \tfrac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s^* - \mu_k\right]\right)\right) = c. \qquad (9)$$

With the normal distribution, which is single-peaked, there generically exists either two or zero solutions to equation (9). Provided that c is sufficiently low for this equation to have two solutions, as assumed in the main text, denote by $\underline{s}$ and $\overline{s}$ as respectively the lowest and highest value of s^* that solve equation (9). We have:

$$\underline{s}(t_k) = \underline{s}(0) + \frac{\sigma_k^2 + \sigma_s^2}{\sigma_k^2} t_k, \quad \text{and} \quad \overline{s}(t_k) = \overline{s}(0) + \frac{\sigma_k^2 + \sigma_s^2}{\sigma_k^2} t_k. \qquad (10)$$

■

Proof of Proposition 1:

We have:

$$
\begin{aligned}
\frac{d}{dt_k}\mathbb{P}\left(\tilde{q}_i \geq t_k | a_i\right) &= -\int_{-\infty}^{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \frac{1}{\sigma_q}\Phi'\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\phi(s_i|a_i)ds_i \\
&\quad -\int_{\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k} \frac{1}{\sigma_q}\Phi'\left(\frac{1}{\sigma_q}\left(t_k - a_i - \overline{e}\right)\right)\phi(s_i|a_i)ds_i \\
&\quad -\int_{\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k}^{\infty} \frac{1}{\sigma_q}\Phi'\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right)\phi(s_i|a_i)ds_i \\
&\quad +\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}\left[\Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right) - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i - \overline{e}\right)\right)\right] \\
&\quad \times\left[\varphi\left(\frac{1}{\sigma_s}\left(\overline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right) - \varphi\left(\frac{1}{\sigma_s}\left(\underline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right)\right],
\end{aligned} \tag{11}
$$

where φ is the PDF of the standard normal distribution. We study the sign of the last two lines on the RHS of equation (11). With $\Phi' > 0$, we have $\Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right) - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i - \overline{e}\right)\right) > 0$. Therefore, the incentive effect is positive if and only if

$$
\varphi\left(\frac{1}{\sigma_s}\left(\overline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right) \geq \varphi\left(\frac{1}{\sigma_s}\left(\underline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right). \tag{12}
$$

Therefore, for given $\underline{s}(0) < \overline{s}(0)$, the condition in (12) holds if and only if

$$
\frac{1}{2}\left[\underline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k + \overline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k\right] \leq a_i \quad \Leftrightarrow \quad t_k \leq \frac{\sigma_k^2}{\sigma_k^2+\sigma_s^2}\left[a_i + \frac{\underline{s}(0)+\overline{s}(0)}{2}\right] \equiv \hat{t}_k(a_i).
$$

This establishes the first part of the Proposition.

For the second part, note that the level of t_k that maximizes the incentive effect for an individual with ability a_i is given by the first-order condition:

$$
\begin{aligned}
&\frac{\partial}{\partial t_k}\Bigg\{\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}\left[\Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i\right)\right) - \Phi\left(\frac{1}{\sigma_q}\left(t_k - a_i - \overline{e}\right)\right)\right] \\
&\times\left[\varphi\left(\frac{1}{\sigma_s}\left(\overline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right) - \varphi\left(\frac{1}{\sigma_s}\left(\underline{s}(0) + \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k - a_i\right)\right)\right]\Bigg\} = 0
\end{aligned}
$$

$$
\begin{aligned}
\Leftrightarrow\quad & \frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2\sigma_q}\left[\Phi'\left(\frac{1}{\sigma_q}\left(t_k-a_i\right)\right)-\Phi'\left(\frac{1}{\sigma_q}\left(t_k-a_i-\overline{e}\right)\right)\right] \\
& \times\left[\varphi\left(\frac{1}{\sigma_s}\left(\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k-a_i\right)\right)-\varphi\left(\frac{1}{\sigma_s}\left(\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k-a_i\right)\right)\right] \\
& +\left(\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}\right)^2\left[\Phi\left(\frac{1}{\sigma_q}\left(t_k-a_i\right)\right)-\Phi\left(\frac{1}{\sigma_q}\left(t_k-a_i-\overline{e}\right)\right)\right] \\
& \times\left[\varphi'\left(\frac{1}{\sigma_s}\left(\overline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k-a_i\right)\right)-\varphi'\left(\frac{1}{\sigma_s}\left(\underline{s}(0)+\frac{\sigma_k^2+\sigma_s^2}{\sigma_k^2}t_k-a_i\right)\right)\right]=0 \quad (13)
\end{aligned}
$$

The derivative of the LHS in (13) with respect to t_k has the opposite sign as the derivative of the LHS in (13) with respect to a_i. Since the equality in (13) holds at the maximum, applying the implicit function theorem establishes the result. ■

Proof of Proposition 2:

With $\sigma_q = 0$, an individual exceeds the threshold if and only if $a_i + e_i \geq t_k$. Thus, with $\sigma_s = 0$, an individual will exert effort if and only if $a_i \in [t_k - \overline{e}, t_k)$ so that we have $\underline{s} = t_k - \overline{e}$ and $\overline{s} = t_k$. In sum:

$$
\mathbb{P}\left(\tilde{q}_i \geq t_k | a_i\right) \quad = \quad \begin{cases} 1 & \text{if } t_k - a_i - e_i < 0, \\ 0 & \text{if } t_k - a_i - e_i > 0. \end{cases} \tag{14}
$$

Let ψ be the PDF of $\tilde{a}_i$, and assume $t_k < \bar{a}$. There are two cases to consider.

First, if $t_k + \bar{e} < \bar{a}$, the ability of all individuals who optimally exert effort is still below $\bar{a}$. The measure of appointed individuals with ability above $\bar{a}$ is then simply:

$$
\Lambda(\bar{a}) \quad = \quad \int_{\bar{a}}^{\infty} \psi\left(a\right) da \tag{15}
$$

Second, if $t_k + \bar{e} \geq \bar{a}$, individuals with talent $a_i \in (\bar{a} - \bar{e}, t_k)$ exert effort and have an ability over $\bar{a}$. The measure of appointed individuals with ability above $\bar{a}$ is then:

$$
\Lambda(\bar{a}) \quad = \quad \int_{\bar{a}-\bar{e}}^{t_k} \psi\left(a\right) da + \int_{\bar{a}}^{\infty} \psi\left(a\right) da. \tag{16}
$$

In sum, for a given $t_k < \bar{a}$, the effect of a change in the threshold t_k on $\Lambda(\bar{a})$ is:

$$
\frac{d}{dt_k}\Lambda(\bar{a}) = \begin{cases} \psi(t_k) & \text{if } \bar{a} \in (t_k, t_k + \overline{e}], \\ 0 & \text{if } \bar{a} > t_k + \bar{e}. \end{cases} \tag{17}
$$

That is, for $t_k < \bar{a}$, we have $\frac{d}{dt_k}\Lambda(\bar{a}) \geq 0$. ■

Proof of Remark 1:

Given a signal s_i and finite constants μ_k and σ_k, the incentive constraint in equation (8) rewrites as:

$$g(t_k, \sigma_s) \equiv \Phi\left(\Sigma^{-1}\left(t_k - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) - \Phi\left(\Sigma^{-1}\left(t_k - \overline{e} - \mu_k - \frac{\sigma_k^2}{\sigma_k^2 + \sigma_s^2}\left[s_i - \mu_k\right]\right)\right) \geq c.$$

We have:

$$\lim_{\sigma_s \to \infty} g(t_k, \sigma_s) = \Phi\left(\Sigma^{-1}\left(t_k - \mu_k\right)\right) - \Phi\left(\Sigma^{-1}\left(t_k - \overline{e} - \mu_k\right)\right)$$

This expression tends to zero as $t_k \to \infty$, so that the incentive constraint is not satisfied for any s_i. ■

about ECGI

The European Corporate Governance Institute has been established to improve *corporate governance through fostering independent scientific research and related activities.*

The ECGI will produce and disseminate high quality research while remaining close to the concerns and interests of corporate, financial and public policy makers. It will draw on the expertise of scholars from numerous countries and bring together a critical mass of expertise and interest to bear on this important subject.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.

Electronic Access to the Working Paper Series

The full set of ECGI working papers can be accessed through the Institute's Web-site (www.ecgi.global/content/working-papers) or SSRN:

Finance Paper Series	http://www.ssrn.com/link/ECGI-Fin.html
Law Paper Series	http://www.ssrn.com/link/ECGI-Law.html