

Selecting the Best: The Persistent Effects of Luck*

Mikhail Drugov,[†] Margaret Meyer[‡] and Marc Möller[§]

July 10, 2024

Abstract

We analyze a model of organizational learning where agents' performance reflects time-invariant unobservable ability, privately-chosen effort, and noise. Our main result is that, even when performance is almost entirely random, maximizing the probability of identifying the best agent (“selective efficiency”) requires biasing final selection in favor of early winners. Making luck persistent, e.g. through fast-tracks, is thus rationalized by the pursuit of selective efficiency. Agents' strategic efforts amplify the persistence of luck. Organizational learning also affects the persistence of initial advantages stemming from identity. Identity-dependent biases, e.g. gender-specific mentoring, create incentives that make selection both more efficient and more equitable.

Keywords: Organizational learning, Incentives, Selective efficiency, Discrimination, Inequality.

JEL classification: D21, D82, D83, M51.

*For valuable comments, we thank Ghazala Azmat, Ángel Hernando-Veciana, Paul Klemperer, Igor Letina, Simon Loertscher, Matthew Mitchell, Inés Moreno de Barreda, Michele Pellizzari, Dmitry Ryvkin, Larry Samuelson, Marta Troya Martinez, and participants in various seminars, workshops, and conferences. All errors are our own.

[†]New Economic School, Universitat Autònoma de Barcelona and Barcelona School of Economics (BSE), and CEPR, mdrugov@nes.ru.

[‡]Nuffield College and Department of Economics, Oxford University, and CEPR, margaret.meyer@nuffield.ox.ac.uk.

[§]University of Bern, marc.moeller@unibe.ch.

1 Introduction

Sometimes an individual’s success is explained, or even discredited, as resulting from an initial stroke of good luck. Frank (2016) documents a multitude of careers of over-achievers, ranging from the arts to business, that were kick-started by fortunate circumstances or events. Such narratives raise the question: To what extent do economic institutions or organizational practices *amplify* the role of luck by making its effects long-lasting?

A common argument, across different social sciences, is that resources, training, mentoring or, more generally, *biases* granted to early strong performers increase the likelihood that an initial stroke of luck translates into a final economic advantage. For example, academic tracking in schools (Gamoran and Mare, 1989) and professional fast-tracks in firms or public agencies (Rosenbaum, 1979; Forbes, 1987; Baker et al., 1994) magnify the importance of early performance for final success.¹ As a consequence of such policies, chance events such as graduating during a recession or being the oldest child in class can have long-lasting effects on both labor market outcomes (Oreopoulos et al., 2012) and educational achievements (Bedard and Dhuey, 2006). Sociologists refer to such phenomena using the term *cumulative advantage* (Merton, 1968, DiPrete and Eirich, 2006) and argue that performance differentials, such as the superior publication records of scientists at elite universities, can be largely explained by accumulated resource advantages rather than inherent differences in talent (Zhang et al., 2022).²

If initial success can be attributed at least in part to *merit*, commonly defined as a combination of ability and effort (Sen, 2000), the use of such biases could be rationalized as improving *selective efficiency*, i.e. the allocation of resources to the most productive individuals. However, in the limit where noise swamps merit in the determination of outcomes, using such biases merely makes luck persistent, by inducing final outcomes to depend on early performance that is almost entirely random. In this paper we demonstrate that, while seemingly at odds with meritocratic principles, making luck persistent in this way is a necessary consequence of an organization’s pursuit of the goal of “selecting the

¹Singapore’s Public Service Leadership Programme boosts the public service careers of the most accomplished college graduates through designated job assignments and leadership workshops (<https://www.psd.gov.sg/leadership/public-service-leadership-careers>).

²In a field experiment on NSF grant proposals Cole et al. (1981) concludes that “the fate of a particular grant application is roughly half determined by the characteristics of the proposal and the principal investigator, and about half by apparently random elements which might be characterized as the luck of the reviewer draw.”

best” in very noisy environments.³ We find that agents’ strategic effort choices, instead of impeding learning, actually improve it but, somewhat paradoxically, lead to a higher persistence of luck. Moreover, when some agents possess initial advantages—“societal luck”—coming from their identity, e.g. race or gender, we show that identity-dependent biases, such as gender-specific mentoring, create incentives that not only make final selection more efficient but also reduce the persistent effect of these initial advantages.

By explaining how institutions shape the role of luck for individual success, our theory helps to illuminate the mechanisms behind economic inequality. This is important because inequality appears to be tolerated when based on merit but not when based on luck (Konow, 2000; Fong, 2001; Cappelen et al., 2007; Cappelen et al., 2013). Stronger beliefs in the relevance of luck increase a country’s social spending (Alesina and Angeletos, 2005) and its citizens’ willingness to implement redistributive policies (Almås et al., 2020). They also affect what recent critics of meritocracy have denoted as the *social divide* between the “deserving” and the “undeserving” (Sandel, 2020). To the extent that political polarization is driven by group identification (Duclos et al., 2004), beliefs about the role of luck for success may influence political outcomes. This is especially relevant when beliefs determine the choice between a low-redistribution “American” equilibrium emphasizing the role of merit and a high-redistribution “European” equilibrium acknowledging the role of luck (Benabou and Tirole, 2006; Alesina et al., 2018).⁴

In Section 2, we present a stylized model of a two-agent, two-stage selection process in which individual performance, at each stage, is the sum of an agent’s time-invariant unobservable ability, privately-chosen effort, and a transitory shock. Agents are ex ante identical to the organization but may share private information about their relative abilities. The organization observes only the ordinal ranking of performances at each stage and attempts to select the more able agent.⁵ Agents choose efforts to maximize the probability of being selected, minus their effort costs.⁶ The organization’s optimal selection rule augments the second-stage performance of the agent who performed better in the

³The term “meritocracy” originates from Young’s (1958) apocalyptic vision of a society in which “merit” serves as the sole determinant of power and wealth. In spite of a debate over what constitutes merit, modern democracies claim to adopt merit as a basis for their allocation of resources and decision-making power, although the validity of this claim has been disputed (Piketty, 2014).

⁴Experimental studies on redistribution find that U.S. subjects implement Gini-coefficients 0.2 points lower when incomes are based on luck than when incomes are based on merit, which would be sufficient to bring down U.S. inequality to European levels (Lefgren et al., 2016).

⁵Ordinal performance measurement arises naturally when performance is hard to quantify. Lazear (2000) documents that for managers, piece rates are employed ten times less frequently than for operatives, and attributes this difference to the absence of a cardinal measure of managerial performance.

⁶Lazear and Rosen (1981) argue that competition for promotions can provide workers inside firms with strong incentives to exert effort and may thus substitute for incentive schemes that rely on cardinal performance measurement when performance is hard to quantify.

first stage with an additive bias and selects the agent who performs better in the second stage. Our main focus of interest is the persistence of early success, i.e. the probability that, in equilibrium, the agent with the better initial performance is selected in the final stage.

We start our analysis in Section 3 by considering the case where agents are as uninformed about their relative abilities as the organization. We first show that effort choices will be identical across agents in *both* stages, in spite of the asymmetries induced by learning from the first-stage performance and the application of the second-stage bias. In the absence of private information, effort choices thus have no effect on selective efficiency, implemented bias, or persistence. Our main result shows that in the limit as noise swamps ability differences as a determinant of performance, equilibrium bias converges to a strictly positive value. In other words, even when ability differences have only negligible impact on performance, equilibrium bias makes first-stage winners considerably more likely to be selected than first-stage losers: *Luck is made persistent*. This shows that the persistence of luck illustrated by our motivating examples need not reflect the use of too much or the wrong kind of bias, but can be understood as a byproduct of the pursuit of selective efficiency.

To provide further insight into the implications of the pursuit of selective efficiency for the persistence of luck, we consider an alternative setting where performance information is *cardinal* rather than ordinal, so bias can condition on the first-stage *margin of victory*. We show that for noise distributions that are normal, or thinner-tailed, organizations will make luck more persistent when individual performance can be ranked but not quantified. Furthermore, the greater persistence of luck under ordinal evaluation is equivalent to greater *front-loading* of the dynamic selection process than in the cardinal case, in that first-stage noise is given a relatively more important role than second-stage noise. As ordinal performance measurement is more prevalent towards the top of an organization's hierarchy, our theory thus highlights the importance of initial luck for selection into positions where selection is most consequential.

In Section 4, we consider the case where agents have some, possibly imperfect, information about their relative abilities. Because, in our setting, effort acts as a substitute for ability in increasing performance, strategic behavior might be expected to decrease the informativeness of the agents' first-stage ranking, thereby reducing or even eliminating the need to make luck persistent through the application of bias. We show that, contrary to this intuition, *informed agents' strategic behavior amplifies the persistence of luck*, because the agent more likely to be better exerts a larger first-stage effort than his rival, thereby reinforcing the agents' ability differential on average. This result resonates

well with the prominent role of biased selection—in the form of fast-tracking and high-potential programs—for careers such as management consulting where collaboration in small, close-knit teams allows workers to learn about their relative abilities.

Finally, in Section 5, we extend our model to allow for a different type of luck—“societal luck”—by assuming that one agent (randomly selected) possesses an “identity” (e.g. race, ethnicity, gender, socio-economic background) that gives him a *transitory* advantage over his rival. For instance, women face disadvantages during the early stages of their careers, e.g. in school grading (Lavy and Megalokonomou, 2024) or assessments of their management potential (Benson et al., 2023). Investigating the mechanisms that propagate such disadvantages, by making their effects long-lasting, therefore ranks high on the agenda of the literature on cumulative advantage, both in economics (Blank, 2005) and sociology (DiPrete and Eirich, 2006).

When we focused on early career luck, we defined persistence as the probability that the agent with better first-stage performance is selected in the final stage. In Section 5, we define the persistence of societal luck as the probability that the agent who receives the transitory, exogenous advantage in the first stage is ultimately selected. We analyze how this persistence is determined by the interaction between agents’ strategic effort choices and the organization’s choice of second-stage bias.

We highlight that an important factor influencing the persistence of societal luck is whether or not organizational selection can condition on whether early success was achieved with or without an exogenous advantage, i.e. whether the second-stage bias can depend on agents’ identity. We show that, if bias must be identity-independent, then organizational learning *always* makes societal luck persistent. In contrast, allowing bias to condition on agents’ identity not only increases the organization’s selective efficiency but also reduces the persistence of societal luck, potentially altogether eliminating its long-term consequences. An important insight is that agents’ strategic effort responses will amplify these beneficial effects, because identity-dependent selection creates incentives that help to mitigate transitory (dis)advantages. These results can be interpreted as showing that affirmative action, e.g. in the form of gender-specific fast-tracks, can yield gains in both efficiency and equity, and more so when workers’ responses to such policies are considered.⁷

⁷Gender-specific fast-tracking exists in both public and private organizations in the form of mentoring programs, such as the United Nations’ Mentorship Program for Emerging Young Women Leaders or the FeMale Talent Program at Accenture. See also Azmat and Boring (2020) for a recent survey on the performance of various gender-based policies in firms.

Related literature Our paper contributes to the literature on organizational learning. Driven by rich evidence about the functioning of internal labor markets (Baker et al., 1994), the seminal studies by Farber and Gibbons (1996), Gibbons and Waldman (1999, 2006), and Altonji and Pierret (2001) have identified firms’ learning about workers’ productivity as a key factor explaining wage and promotion dynamics. A robust empirical finding is that early raises, either in wages or in position, increase the probability of later promotions. Whether this correlation is caused by workers’ inherent productivity differentials or by a “fast-track effect” is controversial, with U.S. evidence in favor of the former (Belzil and Bognanno, 2008) and Japanese evidence pointing towards the latter (Ariga et al., 1999). In the seminal models, serial correlation of promotion rates arises from workers’ time-invariant ability differences or human capital accumulation. Our analysis shows that even when ability differences become negligible and human capital is constant, serial correlation can be explained by the non-vanishing optimality of fast-tracking (bias). The special relevance of early performance for careers is underlined by Lange’s (2007) finding that “employers learn fast”.⁸ Pastorino (2024) supports this view by documenting firms’ tendency to assign newly employed managers to tasks that are particularly informative about their abilities. According to our theory, such task assignments increase the persistence of early career luck even further because they increase the size of the optimal bias. Pastorino’s structural estimates provide strong evidence that learning, in addition to human capital accumulation, has a sizeable impact on career outcomes.

Our theory identifies a mechanism—selection with the help of biases—that augments the relevance of initial performance for final success. Other mechanisms with similar effects exist in recent literature investigating the detailed process through which employers learn about workers’ productivity.⁹ For instance, when organizational learning is viewed as a *bandit problem* (Bergemann and Välimäki, 2008), negative experiences can terminate an employer’s hiring from a group of potential employees or her task assignment to a specific worker. This can lead to long-run disadvantages for groups whose productivity is relatively undiscovered, e.g. minorities (Bartoš et al., 2016; Lepage, 2024), and to persistent discrimination for subgroups of workers in careers which are exposed to “bad news”, e.g. surgeons (Bardhi et al., 2023; Durandard, 2023). Alternatively, when employers rely on the evaluations a worker obtained from previous employers, beliefs about other em-

⁸Using Armed Forces Qualification Test scores as measures of unobserved ability, Lange (2007) finds that it takes only 3 years for employers’ expectation error about workers’ productivity to decline by one half. Similarly, Lluís (2005) finds evidence that employer learning affects mobility between upper and executive levels of German firms but only for workers below 35 years of age. For more experienced workers learning is found to continue to matter when workers differ in how their productivity evolves over time (Kahn and Lange, 2014).

⁹See Sections 3 and 4 of Onuchic (2023) for an excellent survey of this literature.

employers’ information or preferences start to play a role. Focusing on this *social learning* component, [Bohren et al. \(2019\)](#) show that discrimination will be persistent, i.e. negative discrimination will never be reversed, unless it is belief-based and some employers have misspecified priors.¹⁰

A distinguishing feature of our theory is to consider organizational learning as a *strategic interaction*—between a principal choosing bias for selection and agents exerting efforts to become selected. More specifically, our analysis builds on the “pure” organizational learning model of [Meyer \(1991\)](#) but introduces agents’ strategic choices of costly efforts. Decomposing merit into a non-strategic part (“ability”) and a strategic part (“effort”) sheds light on an ongoing controversy over what constitutes merit ([Sen, 2000](#)), by revealing how these two components interact to shape the role of luck for economic outcomes. Notably, the broader notion of merit that makes agents “responsible” for their performance induces organizational selection to assign an even greater role to luck, but only if agents are informed about their relative abilities.

Our emphasis on modeling organizational selection as a strategic interaction is shared by the literature on *equilibrium statistical discrimination* (e.g. [Lundberg and Startz, 1983](#); [Coate and Loury, 1993](#); [Moro and Norman, 2004](#); [Fosgerau et al., 2023](#)). In this literature, workers’ incentives to invest in human capital are a key driver of the result that discrimination can be a self-fulfilling prophecy, and an important insight is that affirmative action—in the form of hiring quotas—can have the perverse effect of reducing investments and wages for members of disadvantaged groups. Our analysis in Section 5, on the contrary, shows that workers’ opportunity to compensate for early, identity-based disadvantages through on-the-job effort constitutes an important channel through which affirmative action—in the form of identity-dependent biases—improves *both* equality *and* selective efficiency.

One component of our analysis in Section 5 of the persistence of societal luck is the recognition that good performance in the face of a disadvantage is particularly informative about an agent’s high ability. This feature is also present in the theoretical analyses of [Meyer \(1991\)](#) and [Fryer \(2007\)](#) and is documented in a field experiment by [Bohren et al. \(2019\)](#). Our analysis also links to the focus in [Sethi and Somanathan \(2023\)](#) and [Bohren et al. \(2023\)](#) on systemic discrimination. In particular, [Sethi and Somanathan \(2023\)](#) argue that expanding the representation of disadvantaged groups in hiring, e.g. through exploration-prone algorithms ([Li et al., 2024](#)), can be *necessary* for selective efficiency. Unlike [Sethi and Somanathan \(2023\)](#) and [Bohren et al. \(2023\)](#), however, we highlight that

¹⁰[Gill and Prowse \(2014\)](#) experimentally document a different mechanism by which initial performance influences final success, namely, the psychological impact of early wins and losses on subsequent effort choices. They find that this impact differs systematically between men and women.

the effort responses to identity-based policies amplify their beneficial effects with respect to both learning and persistence.

2 Model

We consider an organization consisting of a risk-neutral principal and two agents $i \in \{A, B\}$ with heterogeneous abilities a_i . The difference in abilities or *heterogeneity* is given by $h > 0$, i.e. $\Delta a \equiv a_A - a_B \in \{-h, h\}$. The principal observes the agents' relative performance during two stages, $t \in \{1, 2\}$. After the second stage, the principal needs to select one of the agents for a higher-level task whose payoff to the principal is increasing in the selected agent's ability. The principal's goal is thus simple: to select the more able agent.

Agent i 's performance at stage t , $x_{i,t} \in \mathfrak{R}$, is the sum of three elements: the agent's time-invariant ability a_i , multiplied by a stage-specific weight $\lambda_t > 0$; the agent's private choice of effort $e_{i,t} \geq 0$; and a time-varying random component $\epsilon_{i,t} \in \mathfrak{R}$. That is,

$$x_{i,t} \equiv \lambda_t a_i + e_{i,t} + \epsilon_{i,t}.$$

Variation in λ_t across stages can reflect differences in the impact of ability on performance. This is especially relevant when agents' tasks change over time.

Information and choices The principal and the agents share a common prior, $q^0 \equiv \mathbb{P}(\Delta a = h) \geq \frac{1}{2}$, but for the principal, agents A and B are indistinguishable. If $q^0 = \frac{1}{2}$, the agents are as uninformed as the principal, while if $q^0 > \frac{1}{2}$, the agents have superior information about their relative abilities, with both agents believing that agent A is more likely to be better.¹¹

The principal can observe only the ranking of the two agents' performances after the first stage. In the second stage, the principal may costlessly and publicly assign a bias $\beta \in \mathfrak{R}$ to the winner of the first stage. If $\beta > 0$, the bias increments the winner's second-stage performance, and we say that the bias "favors" the first-stage winner, whereas if $\beta < 0$, it reduces his second-stage performance. Having won the first stage, agent i is then identified as the winner of the second stage if $x_{i,2} + \beta > x_{j,2}$.¹²

¹¹Virtually all the employer learning models reviewed in the Introduction analyze only the case where workers are ignorant about their own ability, and hence correspond to the case $q^0 = \frac{1}{2}$. Section 4 contrasts the results for $q_0 > \frac{1}{2}$ with those in Section 3 for $q_0 = \frac{1}{2}$.

¹²In Section 3.2 we analyze a cardinal setting where the principal observes the difference in the first-stage performances and can condition the choice of bias on this difference.

The principal chooses the size of bias β and the selection rule to maximize *selective efficiency*, $S(\beta; h)$, defined as the probability that the more able agent is selected. The agents exert privately-observed efforts $e_{i,t}$ in each stage to maximize the probability of being selected minus the effort costs. The value of being selected is the same for both agents and is normalized to 1. The cost-of-effort functions $C_t(e_{i,t})$, $t = 1, 2$, are strictly increasing and convex. Thus, effort costs are identical across agents but may differ across stages.

Noise The distribution of the difference in the individual noise terms, $\Delta\epsilon_t \equiv \epsilon_{A,t} - \epsilon_{B,t}$, is a key primitive in our model because outcomes depend only on performance *differentials*. We assume that $\Delta\epsilon_t$ are identically and independently distributed across stages and denote the corresponding support by $[-z, z]$ (where z may be infinite), the cumulative distribution function by G , and its density by g . We make the following distributional assumptions:

Assumption 1 (i) g is symmetric about 0; (ii) g is strictly log-concave; (iii) g is differentiable on $(-z, z)$; (iv) $\lim_{y \rightarrow z} L(y) = \infty$, where

$$L(y) \equiv -\frac{g'(y)}{g(y)}.$$

The symmetry of g captures the idea that the only source of heterogeneity across agents is the difference in their abilities; it is a weaker assumption than individual shocks, $\epsilon_{i,t}$, being i.i.d. across agents. Log-concavity of g is equivalent to the monotone likelihood ratio property in our setting; it guarantees that, in either stage, a larger performance differential $\Delta x_t \equiv x_{A,t} - x_{B,t}$ implies a higher likelihood that A 's ability exceeds B 's. It also implies that L is increasing. Strict log-concavity makes all the implications strict. Together with the remaining two assumptions, it ensures that the principal's problem is well-behaved.

Timing In the beginning of the first stage, the agents choose efforts. Then, the noise is realized, and both the principal and the agents observe who has higher first-stage performance. In the beginning of the second stage, the principal chooses the level of bias. Then the agents exert efforts. The noise is realized, and both the principal and the agents observe who has a higher second-stage performance. The principal then selects one of the agents. Note that the principal chooses the bias *after* the first stage rather than committing to it in the beginning.¹³

¹³Assuming instead that the principal can commit, before agents choose efforts, to a value of bias (in the case of ordinal performance evaluation analyzed in Section 3.1) or to a schedule of bias (in case of

Equilibrium The solution concept is perfect Bayesian equilibrium (PBE). In a PBE, (i) the effort choice by each agent at each stage is optimal for him given his conjectures about the effort choices of the other agent and the bias and the selection rule set by the principal; (ii) the bias and the selection rule are optimal for the principal given her conjectures about the agents' efforts; and (iii) the conjectures of both agents and the principal are correct.

It is easy to confirm that when the principal chooses the bias optimally, the optimal selection rule is to select the winner of the second stage.

Persistence Our main focus of interest is the persistence of outcomes induced by the interaction between the principal's pursuit of selective efficiency and the agents' desire to be selected. We define *persistence* as the equilibrium probability that the winner of the first stage is selected after the second stage.

It is important to note that the key parameter of our model, $h > 0$, which captures the degree of agents' heterogeneity in abilities, also has a broader interpretation as the *ratio* of agents' heterogeneity to the scale of noise.¹⁴ To shed light on the role of initial luck for final outcomes, much of the analysis in Sections 3 and 4 will focus on the setting in which h is very small: Here the scale of noise is large *relative* to the agents' heterogeneity and, as we show, differences in agents' efforts vanish. Note that even in this environment, the selection decision may still be important to the principal, because the selected agent's performance in the higher-level task may be very sensitive to ability.

When, in this limiting environment, persistence turns out to be strictly larger than one-half, we will say that *luck is made persistent*, because the first-stage winner has a greater chance of ultimately being selected than the first-stage loser, despite the fact that the first-stage outcome is almost entirely determined by random factors.

3 Uninformed agents

In this section, we consider the setting where agents are as uninformed as the principal about their relative abilities. We show that, in this case, the agents' ability to influence their performance through the exertion of effort has no impact on selective efficiency, and hence no impact on the principal's choice of bias or on the persistence of early success. This

cardinal performance evaluation analyzed in Section 3.2) would not affect our results unless agents are informed about their relative abilities (see Section 4).

¹⁴To see this, introduce a scaling transformation $\Delta\epsilon_t \rightarrow \sigma\Delta\epsilon_t$, with $\sigma > 1$, which makes the difference in the noise terms more dispersed: The cdf becomes $G(\frac{\Delta\epsilon_t}{\sigma})$, the pdf $\frac{1}{\sigma}g(\frac{\Delta\epsilon_t}{\sigma})$, and the support $[-\sigma z, \sigma z]$. If the underlying heterogeneity in abilities is H , then $G(\frac{\lambda_1 H}{\sigma})$ is the probability that, when the first-stage effort differential is zero, the more able agent wins the first stage. It depends on H and σ only through the *heterogeneity-to-noise ratio* $h \equiv \frac{H}{\sigma}$.

allows us to develop the basic intuition for the connection between these variables, before examining, in Section 4, the effects of informed agents’ strategic behavior. In Section 3.1 we prove our main result, that even as noise swamps ability differences, equilibrium bias converges to a strictly positive value, so luck is made persistent. In Section 3.2 we analyze an alternative case where performance evaluation is cardinal rather than ordinal, allowing bias to depend on the first-stage margin of victory. We show that, under mild conditions on the distribution of noise, the limiting equilibrium bias under ordinal evaluation front-loads the dynamic selection process compared to the limiting equilibrium bias under cardinal information.

3.1 Equilibrium bias and persistence of luck

We start with the following lemma, which shows that, in equilibrium, the efforts of uninformed agents cancel each other in the determination of relative performance. The proof of this lemma and all other proofs are in the Appendix.

Lemma 1 (Identical efforts) *Let $q^0 = \frac{1}{2}$. Then for any anticipated choice of bias β by the principal, there exists a unique pure-strategy equilibrium in efforts. In this equilibrium, agents choose identical efforts, both in the first stage and in the second stage.*

In the second stage, despite the asymmetries due to learning and the use of bias, the marginal benefit of effort is the same for the two agents. This is because the values of winning, the marginal impacts of effort on performance, and the pivotal realizations of $\Delta\epsilon_2$ are all identical for A and B (cf. Lazear and Rosen, 1981). In the first stage, given the symmetry of the agents’ situations, there exists a pair of identical efforts that are best responses to each other. If $\beta \leq 0$, a higher first-stage performance reduces or leaves unchanged the probability of being selected after the second stage, so only zero efforts are mutual best responses. If $\beta > 0$, we show by contradiction that unequal efforts could not be best responses. Specifically, if agent A were to exert more effort than agent B in the first stage, then a first-stage win by B would be a stronger signal of ability than a win by A . Hence, the biased second-stage contest would be more unbalanced following a win by B and would therefore induce lower second-stage effort. But lower second-stage efforts after a win by B would generate stronger first-stage incentives for B than for A , which contradicts the initial assumption.

Given Lemma 1, the equilibrium efforts cancel out in selective efficiency $S(\beta; h)$, the probability with which the more able agent wins the second stage:

$$S(\beta; h) = G(\lambda_1 h)G(\lambda_2 h + \beta) + [1 - G(\lambda_1 h)]G(\lambda_2 h - \beta). \quad (1)$$

The first term in the sum is the probability that the more able agent wins the first stage and then wins the second stage with bias β in his favor. The second term is the probability that the more able agent loses the first stage but then wins the second stage despite being disadvantaged by the bias.

The equilibrium value of bias is then found by maximizing (1) with respect to β , which yields the following first-order condition:¹⁵

$$\frac{G(\lambda_1 h)}{1 - G(\lambda_1 h)} = \frac{g(\lambda_2 h - \beta)}{g(\lambda_2 h + \beta)}. \quad (2)$$

The ratio on the left-hand side is the relative likelihood that a first-stage win is achieved by the more able agent compared to the less able one. A victory in the first stage is a stronger signal about relative ability when this likelihood ratio is higher. The term on the right-hand side is also a likelihood ratio: It is the relative likelihood that a second-stage draw when agent j is disadvantaged by bias β , i.e. $x_{j,2} - \beta = x_{i,2}$, is achieved when j is the more able agent compared to when j is the less able one. Equation (2) shows that equilibrium bias strikes a balance between the informativeness of the ordinal first-stage ranking—an unbiased win—and the informativeness of the marginal second-stage outcome—a draw achieved despite being handicapped by bias. Equilibrium bias is such that, if the principal were to observe a draw in stage two, she would be indifferent about which agent to select.

Note that, for $\beta = 0$, the right-hand side of (2) is equal to one and hence strictly smaller than the left-hand side. This is because, for $\beta = 0$, a second-stage draw is uninformative about the agents' abilities. Moreover, given the strict log-concavity of g , as the size of the bias disadvantaging the first-stage loser increases, a second-stage draw becomes a strictly stronger signal about that agent's relative ability. It thus follows from Assumption 1 that the first-order condition (2) has a unique solution, $\beta^*(h) > 0$, which maximizes selective efficiency. Moreover, since the left-hand (right-hand) side of (2) is increasing in λ_1 (λ_2), which measures the sensitivity of the first-stage (second-stage) performance to ability, $\beta^*(h)$ is increasing in λ_1 and decreasing in λ_2 .

While these arguments establish that a positive bias will emerge in equilibrium for any level $h > 0$ of heterogeneity in abilities, they are not sufficient to determine what happens in the limit as $h \rightarrow 0$. Does equilibrium bias converge to zero? The following proposition characterizes the limiting value of the equilibrium bias as the scale of the noise swamps the heterogeneity in abilities.

¹⁵Given that the agents choose identical first-stage efforts (Lemma 1), this first-order condition matches the one in the pure organizational learning model of Meyer (1991).

Proposition 1 (Equilibrium bias) *Let $q^0 = \frac{1}{2}$. The principal's equilibrium choice of bias, $\beta^*(h)$, is strictly positive, even in the limit as noise swamps agents' ability differences. More specifically, $\beta_0^* \equiv \lim_{h \rightarrow 0} \beta^*(h) > 0$ is given by the unique solution of the equation*

$$2\lambda_1 g(0) = \lambda_2 L(\beta_0^*). \quad (3)$$

At first sight, the fact that equilibrium bias remains strictly positive, even in the limit, may seem counterintuitive, because when h tends to zero, a first-stage win becomes completely uninformative about relative abilities. However, this reasoning neglects the fact that, as h tends to zero, a second-stage draw also becomes uninformative, for any level of bias. Formally, as h tends to zero, both sides of equation (2) approach one. Proposition 1 thus characterizes equilibrium bias in this limit by equating the *rates* at which the informativeness of each stage tends to zero as h gets small. Since L is a strictly increasing function, $L(0) = 0$, and the left-hand side of (3) is positive, the limiting value of bias must be positive. More intuitively, observe that, when bias is zero, achieving a second-stage draw is uninformative about relative abilities *for any* ratio h of heterogeneity to noise, whereas the informativeness of a first-stage win rises with h . Thus, a strictly positive bias emerges in the limit because, unless first-stage losers are disadvantaged relative to first-stage winners even when ability differences are negligible, the informativeness of a second-stage draw cannot keep up with the informativeness of a first-stage win when ability differences start to matter.

An alternative interpretation of the limiting value of equilibrium bias is illustrated in Figure 1. In the limit as $h \rightarrow 0$, bias is chosen to maximize not the *level* of selective efficiency—since selective efficiency becomes independent of bias in the limit—but the *rate* at which selective efficiency increases with the agents' heterogeneity. In the limit, equilibrium bias thus maximizes the potential gains to selective efficiency from a marginal increase in agents' heterogeneity; were bias set to zero, these gains would not be fully realized.

Though the logic behind the equilibrium level of bias is clear in the limit, the dependence of $\beta^*(h)$ on the heterogeneity-to-noise ratio for $h > 0$ can be complex. This is because an increment in h increases both sides of equation (2): It raises *both* the informativeness of a first-stage win *and*—by log-concavity of g —the informativeness of a second-stage draw, for any given level of bias. The complex dependence of $\beta^*(h)$ on h is illustrated in Figure 2. The left panel plots the density functions for the family of exponential power distributions with mean zero and shape parameter $\eta > 1$.¹⁶ The right panel

¹⁶ These density functions are given by $g(\Delta\epsilon_t; \eta) = \frac{\eta}{2\Gamma(\frac{1}{\eta})} \exp(-|\Delta\epsilon_t|^\eta)$, and for all $\eta > 1$, they satisfy




Figure 1: **Selective efficiency.** The figure depicts selective efficiency S as a function of agents' heterogeneity h for different values of bias. $\beta_0^* > 0$ maximizes the slope of S at $h = 0$.




Figure 2: **Example distributions of noise and equilibrium bias.** The left panel depicts the density functions when noise follows an exponential power distributions with mean zero and shape parameter $\eta \in \{1.5, 2, 3, 4, 7\}$. The right panel plots the corresponding equilibrium bias as a function of h when $\lambda_1 = \lambda_2 = 1$.

Assumption 1. For $\eta = 2$, $g(\Delta\epsilon_t; \eta)$ is a normal distribution with variance $\frac{1}{2}$; as $\eta \rightarrow \infty$, $g(\Delta\epsilon_t; \eta)$ approaches a uniform distribution with support $[-1, 1]$; and as $\eta \rightarrow 1$, $g(\Delta\epsilon_t; \eta)$ approaches a Laplace distribution with scale parameter 1. At $\eta = 1$, Assumption 1 is violated because the Laplace density is not differentiable at 0 and is not strictly log-concave.

in Figure 2 plots the equilibrium bias $\beta^*(h)$ as a function of h , for $\lambda_1 = \lambda_2 = 1$. Despite the many possibilities illustrated, we see that, as shown by Proposition 1, equilibrium bias remains positive even as h gets small for all members of the family.

Persistence

Our results have implications for our understanding of the relevance of luck for the determination of economic outcomes. According to meritocratic principles, the allocation of resources and decision-making power should be attributable to *merit*—a combination of ability and effort—rather than luck. In light of this principle, it is important to ask how institutions and organizational practices shape the dynamic relationship between performance and outcomes. A straightforward but important implication of the introduction of bias favoring the first-stage winner is that it raises the correlation between initial success and final selection. To see this, define the *persistence* of the selection process as the equilibrium probability that the first-stage winner is selected after the second stage. Given Lemma 1, persistence is independent of efforts and is given by:

$$P(\beta^*(h); h) = G(\lambda_1 h)G(\lambda_2 h + \beta^*(h)) + [1 - G(\lambda_1 h)][1 - G(\lambda_2 h - \beta^*(h))]. \quad (4)$$

Of course, even in the absence of bias, initial success and final selection are positively correlated, and hence $P(0; h) > \frac{1}{2}$, because the outcomes of both stages are affected by the time-invariant ability difference $h > 0$. However, in the limit as $h \rightarrow 0$, this correlation would vanish, and hence persistence would approach $\frac{1}{2}$, *unless* it were induced through the use of bias. That is, defining $P_0^* \equiv \lim_{h \rightarrow 0} P(\beta^*(h); h)$, we have from (4) that

$$P_0^* = G(\beta_0^*) \quad \text{and} \quad P_0^* > \frac{1}{2} \iff \beta_0^* > 0. \quad (5)$$

Hence, a direct implication of Proposition 1 is that *luck is made persistent*: $P_0^* > \frac{1}{2}$. Also note that (3), coupled with the strict monotonicity of L , implies that β_0^* , and hence P_0^* , is increasing in the ratio λ_1/λ_2 , which measures the relative sensitivity to ability of first-stage compared to second-stage performance. This is true even though in the limit, ability has only a negligible impact on performance. These important implications are stated in the next corollary.

Corollary 1 (Persistence of luck) *Let $q^0 = \frac{1}{2}$. When bias is set to maximize selective efficiency, luck is made persistent, i.e. $P_0^* > \frac{1}{2}$, and even more so when early performance is relatively more sensitive to ability, i.e. P_0^* is strictly increasing in $\frac{\lambda_1}{\lambda_2}$.*

Recent work by [Pastorino \(2024\)](#) shows that firms tend to allocate to newly-hired workers those tasks that are relatively more informative about their abilities; our results show that this pattern of task allocation enhances the persistence of luck.

In addition, because persistence in (4) is increasing both in the bias and in heterogeneity, the fact that, as shown by Figure 2, equilibrium bias $\beta^*(h)$ can be decreasing suggests that, overall, persistence could also be decreasing in h . This possibility is confirmed by Figure 3, where for $\eta = 7$ and relatively small h , equilibrium persistence falls as the difference in agents' abilities rises. This means that the use of bias for selection can make final success *less* correlated with initial performance in settings where performance differentials are *more* attributable to ability differences.




Figure 3: **Equilibrium persistence.** The figure plots the probability $P(\beta^*(h); h)$ that the first-stage winner is ultimately selected, as a function of the ratio h of agents' heterogeneity to noise, when noise follows an exponential power distribution with mean 0 and shape parameter $\eta \in \{1.5, 2, 3, 4, 7\}$ and $\lambda_1 = \lambda_2 = 1$.

To conclude, our analysis shows that two apparent violations of meritocratic principles—the persistence of luck and non-monotonicity of equilibrium persistence—can be rationalized by the very fact that organizations aim to allocate resources to the most talented individuals. Thus, neither of these phenomena should automatically be considered an abandonment of meritocratic principles. In fact, Corollary 1 shows that making luck persistent by biasing selection in favor of early strong performers emerges as a *necessary* consequence of the meritocratic pursuit of selective efficiency.

3.2 Cardinal performance evaluation

We now analyze how the use of bias, and its consequences for the persistence of luck, vary with the way in which performance differentials are measured, contrasting the case of *ordinal* information, studied so far, to that of *cardinal* information. Lazear (2018) argues that ordinal performance evaluation is prevalent towards the top of an organization’s hierarchy, given the difficulty of quantifying the performance of increasingly complex tasks. This means that in situations where selection matters most, for both the organization and the agents themselves, ordinal performance measurement may be the most relevant case. However, a comparison with the case where the principal can quantify the agents’ performance differentials helps to highlight the specific contribution of rank-order information to the persistence of luck. It may also help to assess whether luck can be expected to play a more important role for selection into positions with higher ranks.

When all parties can observe the first-stage performance differential $\Delta x_1 = x_{A,1} - x_{B,1}$, and can condition their second-stage actions on it, there exists an equilibrium in which, in both stages, agent A exerts the same effort as agent B .¹⁷ Hence, in this equilibrium, similarly to Section 3.1, efforts do not matter for selective efficiency or for persistence. The probability of a *margin of victory* $|\Delta x_1|$ being achieved by the stronger agent is $g(|\Delta x_1| - \lambda_1 h)$, whereas for the weaker agent the corresponding probability is $g(|\Delta x_1| + \lambda_1 h)$. Given the observed $|\Delta x_1|$, the principal then chooses the bias to maximize

$$S^{card}(\beta, |\Delta x_1|; h) = g(|\Delta x_1| - \lambda_1 h)G(\lambda_2 h + \beta) + g(|\Delta x_1| + \lambda_1 h)G(\lambda_2 h - \beta). \quad (6)$$

Intuitively, a larger margin of victory $|\Delta x_1|$ is a stronger signal about the winner’s ability and thus induces the principal to choose a larger bias $\beta^{card}(|\Delta x_1|, h)$.

Equilibrium bias under cardinal information is particularly transparent when performance in the two stages is equally sensitive to ability, that is, when $\lambda_1 = \lambda_2$. Here it is optimal for the principal to select the agent with the higher *aggregate performance*, $x_{i,1} + x_{i,2}$. This selection rule can be implemented by biasing the second stage in favor of the first-stage winner by exactly $|\Delta x_1|$, the first-stage margin of victory. Hence, in this case, $\beta^{card}(|\Delta x_1|, h) = |\Delta x_1|$, for all $|\Delta x_1|$ and h .

In general, the equilibrium bias when performance evaluation is ordinal, $\beta^*(h)$ given by (2), can be thought of as a form of average of the equilibrium biases $\beta^{card}(|\Delta x_1|, h)$ under cardinal evaluation, as $|\Delta x_1|$ varies. Proposition 2 makes this intuition precise for the limiting case where noise swamps ability. We define $\beta_0^{card}(|\Delta x_1|) \equiv \lim_{h \rightarrow 0} \beta^{card}(|\Delta x_1|, h)$.

¹⁷In close analogy to Lemma 1, agents exerting identical efforts constitutes the *unique* pure-strategy equilibrium if agents anticipate that the principal chooses bias optimally, based on cardinal performance information and given her conjecture about agents’ efforts.

Proposition 2 (Cardinal bias) Let $q^0 = \frac{1}{2}$. When the principal can condition bias on cardinal performance information $|\Delta x_1|$, the following holds as $h \rightarrow 0$:

(i) $\beta_0^{card}(|\Delta x_1|) > 0$ whenever $|\Delta x_1| > 0$, and $\beta_0^{card}(|\Delta x_1|)$ solves

$$L(\beta_0^{card}(|\Delta x_1|)) = \frac{\lambda_1}{\lambda_2} L(|\Delta x_1|). \quad (7)$$

(ii) Cardinal bias and ordinal bias are related according to

$$\mathbb{E}[L(\beta_0^{card}(|\Delta x_1|))] = L(\beta_0^*). \quad (8)$$

Note that when the difference in the agents' noise terms is normally distributed, so L is linear, (8) implies $\mathbb{E}[\beta_0^{card}(|\Delta x_1|)] = \beta_0^*$.

A direct implication of Proposition 2(i), recalling (4) and (5), is that with cardinal performance evaluation, luck is made persistent *on average*, i.e.

$$P_0^{card} \equiv \lim_{h \rightarrow 0} \mathbb{E}[P(\beta^{card}(|\Delta x_1|, h), h)] = \mathbb{E}[G(\beta_0^{card}(|\Delta x_1|))] > \frac{1}{2}. \quad (9)$$

For the special case of $\lambda_1 = \lambda_2$, since the principal selects the agent with the higher aggregate performance $x_{i,1} + x_{i,2}$, we have, for any noise distribution g ,

$$P_0^{card} = \lim_{h \rightarrow 0} \mathbb{P}(\Delta x_1 + \Delta x_2 \geq 0 | \Delta x_1 \geq 0) = \mathbb{P}(\Delta \epsilon_1 + \Delta \epsilon_2 \geq 0 | \Delta \epsilon_1 \geq 0) = \frac{1}{2} + \frac{1}{4} = \frac{3}{4},$$

since $\Delta \epsilon_1$ and $\Delta \epsilon_2$ are i.i.d.

To highlight the specific contribution of ordinal evaluation to the persistence of luck, we compare P_0^{card} with P_0^* in (5), the persistence under ordinal evaluation. Figure 3, which is plotted for $\lambda_1 = \lambda_2 = 1$, shows that P_0^* is larger than $\frac{3}{4}$ for all values of the shape parameter depicted. For $\eta = 2$ ($\Delta \epsilon_t$ normally distributed), this is not surprising, given that $\mathbb{E}[\beta_0^{card}(|\Delta x_1|)] = \beta_0^*$ and given that persistence in (9) is the expectation of a concave function of bias.¹⁸ In fact, a sufficient condition for P_0^* to exceed P_0^{card} is that the function L is convex, since (8) shows that for L convex, the limiting ordinal bias is at least as large as the expected limiting cardinal bias.¹⁹ Distributions for which L is convex are those with densities \tilde{g} that are *thinner-tailed* than the normal distribution, more precisely, those that are more log-concave than the normal in the sense that $\ln \tilde{g}$ is a concave transform of $\ln g$,

¹⁸Concavity of G on the positive domain follows from the log-concavity and symmetry about 0 of g .

¹⁹Convexity of L is not necessary for P_0^* to exceed P_0^{card} . For the exponential power family of distributions in footnote 16, L is convex if and only if $\eta \geq 2$, but for $\lambda_1 = \lambda_2$, the persistence of luck is larger under ordinal than under cardinal evaluation for all $\eta > \sim 1.38$.

for g normal.

The following corollary shows that this insight extends to the case of arbitrary λ_1 and λ_2 .

Corollary 2 (Persistence of Luck: Cardinal versus ordinal evaluation) *Let $q^0 = \frac{1}{2}$ and suppose that the function L is convex.*

(i) *The persistence of luck is greater when performance evaluation is ordinal than when it is cardinal, i.e.*

$$P_0^* = G(\beta_0^*) > \mathbb{E}[G(\beta_0^{card}(|\Delta x_1|))] = P_0^{card}. \quad (10)$$

(ii) *The inequality in (10) is equivalent to the organization assigning greater relative weight to first-stage performance than to second-stage performance when performance evaluation is ordinal than when it is cardinal, as $h \rightarrow 0$, i.e. for all λ_1, λ_2 ,*

$$\frac{\mathbb{P}(\text{select } A | \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{ord.})}{\mathbb{P}(\text{select } A | \Delta\epsilon_1 < 0, \Delta\epsilon_2 > 0, \text{ord.})} > \frac{\mathbb{P}(\text{select } A | \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{card.})}{\mathbb{P}(\text{select } A | \Delta\epsilon_1 < 0, \Delta\epsilon_2 > 0, \text{card.})}. \quad (11)$$

Corollary 2 shows that if, as argued by Lazear (2000), organizations are constrained to use ordinal performance measurement at high ranks because of the difficulty of quantifying performance in complex tasks, luck may have especially persistent effects on selection at the top of the hierarchy.

Corollary 2 also shows that greater persistence of luck under ordinal than cardinal evaluation is equivalent to greater “front-loading” of the dynamic selection process when performance evaluation is constrained to be ordinal. This inflation of the importance of early luck under ordinal evaluation is especially transparent when performance in the two stages is equally sensitive to ability, that is, when $\lambda_1 = \lambda_2$: Whereas under cardinal evaluation luck is weighted equally across stages (the right-hand side of (11) equals one), under ordinal evaluation early luck has a greater impact on selection than later luck (the left-hand side of (11) is greater than one).

4 Informed agents

The model in Section 3 shares with the literature on organizational learning (e.g. Gibbons and Waldman, 2006; Lange, 2007; Pastorino, 2024) the assumption that agents are as uninformed about their relative abilities as the principal. But what if agents have

some common, possibly imperfect, information about their relative abilities? For example, workers might know each other from college or might have shared experiences with previous employers, allowing them to better judge their relative abilities. We capture this by assuming that $q^0 \equiv \mathbb{P}(\Delta a = h) > \frac{1}{2}$, where $\Delta a = a_A - a_B$. While with uninformed agents, Lemma 1 showed that effort choices had no impact on the principal’s learning, because in equilibrium efforts cancelled each other in the determination of relative performance, with informed agents, their efforts may no longer be identical. In this section, we examine how the strategic behavior of informed agents impacts organizational learning and the persistence of luck, in our main setting of ordinal performance evaluation.

Because effort and ability are *substitutes* for agents’ performance, the agent thought less likely to be the more able might use effort to compensate for his ability disadvantage, thereby decreasing the informativeness of early performance about relative abilities, reducing the optimal bias, and weakening our result about the persistence of luck. In fact, we show that, on the contrary, informed agents’ strategic behavior *reinforces* the impact of agents’ ability difference, resulting in luck being made even more persistent than when agents are ignorant of relative abilities. The following lemma represents the crucial step in our argument.

Lemma 2 (Informed agents’ effort differential) *Let $q^0 > \frac{1}{2}$. Then for any anticipated choice of bias $\beta > 0$, agents choose identical efforts in the second stage, but in the first stage, the agent thought more likely to have higher ability exerts a strictly larger effort than his rival.*

The explanation for why the agents choose identical second-stage efforts is the same as for Lemma 1. To understand the sign of the first-stage effort differential

$$\Delta e_1^*(\beta, h, q^0) \equiv e_{A,1}^*(\beta, h, q^0) - e_{B,1}^*(\beta, h, q^0) > 0,$$

note first that because exactly one agent will be selected after the second stage, the “rewards” of winning the first stage arising from the increased probability of being selected are precisely the same for the two agents. However, in contrast to the case where agents are uninformed, the *level* of second-stage effort that agents exert, and hence their effort cost, now depends on which agent wins the first stage. To see this most clearly, suppose for simplicity that $q^0 = 1$, so that agent A is known with certainty to be more able. Recall that the principal is aware of the agents’ superior knowledge but cannot distinguish agent A from agent B , so she must assign the same level of bias whoever wins the first stage. If agent A wins the first stage, then the bias will reinforce the agents’ ability difference, and the pivotal realization $h + \beta$ of noise $\Delta \epsilon_2$ will determine second-stage efforts via

$C'_2(e_{A,2}^*) = g(h + \beta) = C'_2(e_{B,2}^*)$. If, instead, agent A loses the first stage, then bias will mitigate the agents' ability difference, so it is the pivotal realization $h - \beta$ that determines second-stage efforts via $C'_2(e_{A,2}^*) = g(h - \beta) = C'_2(e_{B,2}^*)$. Because $g(h + \beta) < g(h - \beta)$ by unimodality of g (implied by log-concavity), agent A faces lower second-stage effort costs after winning the first stage than after losing, so A has a “cost-saving incentive” to win the first stage. For agent B , the argument is reversed, because bias mitigates agents' heterogeneity when B wins but reinforces it when B loses, so agent B has a “cost-saving disincentive” for first-stage effort.

Lemma 2 shows that, in equilibrium, informed agents' first-stage effort differential on average reinforces the ability difference, thus raising the informativeness of the first-stage outcome. The following result extends Proposition 1 to the case of informed agents, under the additional assumption that effort costs are quadratic.²⁰

Proposition 3 (Bias with informed agents) *Let $q^0 > \frac{1}{2}$ and suppose that $C_t(e_{i,t}) = \frac{c_t}{2}e_{i,t}^2$ for all i, t . In the limit as noise swamps ability differences, equilibrium bias $\beta_0^*(q^0) \equiv \lim_{h \rightarrow 0} \beta^*(h, q^0)$ is unique, strictly positive, and strictly increasing in q^0 , and it solves*

$$2g(0) \left[\lambda_1 + (2q^0 - 1) \frac{\partial \Delta e_1(\beta_0^*(q^0), 0, q^0)}{\partial h} \right] = \lambda_2 L(\beta_0^*(q^0)). \quad (12)$$

Proposition 3 shows that our insights about the optimal use of bias for selection are robust to the introduction of private information on the part of the agents about their relative abilities. In particular, equilibrium bias continues to remain positive in the limit as noise swamps ability differences. Even though in this limit, the first-stage effort differential $\Delta e_1(\beta, h, q^0)$ between the “better” agent A and the “worse” agent B vanishes, nevertheless $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1(\beta, h, q^0)}{\partial h} > 0$, which implies that the informativeness of the first-stage outcome increases as h rises from 0; hence, the left-hand side of (12), just like the left-hand side of (3), is strictly positive, ensuring that equilibrium bias $\beta_0^*(q^0)$ is strictly positive.

Proposition 3 also reveals that the limiting equilibrium bias is strictly increasing in the precision q^0 of the agents' private information and that there are two distinct forces generating this result. First, (12) shows that the larger is q^0 , the greater is the impact of any given $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$, because the effort differential is more likely to be aligned with the ability difference. Second, the larger is q^0 , the larger is $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$ itself, because in the first stage, both A 's “cost-saving incentive” for effort and B 's “cost-saving disincentive” are stronger the better informed the agents are about relative abilities.

This comparative statics result in Proposition 3 provides further insights about the

²⁰The assumption of quadratic costs simplifies the proof that the equilibrium in the limit as $h \rightarrow 0$ is unique, but it is not necessary for this result.

relevance of luck for selection:

Corollary 3 (Persistence amplified) *When agents have private information about their relative abilities, luck is made even more persistent than when agents are uninformed, i.e. $P_0^*(q^0) = G(\beta_0^*(q^0)) > P_0^*$ for all $q^0 > \frac{1}{2}$. Moreover, the persistence of luck is greater the more informed the agents are about their relative abilities, i.e. $P_0^*(q^0)$ is increasing.*

Corollary 3 emphasizes that the more informed the agents are about their relative abilities, the more their strategic behavior amplifies the persistence of luck induced by the organization’s use of bias. The corollary also relates our theory to an ongoing discussion of what constitutes “merit” (Sen, 2000). Inherited talents, acquired abilities, and costly noble acts are all potential sources of merit, endowing their possessor with a justification for receiving decision-making power or economic prosperity. Our theory allows us to distinguish between the case where performance—or merit—is given by the (noisy) sum of an agent’s ability *and* effort, and the case where only ability matters. Section 3 showed that whether or not effort is included in the definition of merit is irrelevant for the outcome of organizational selection when agents are uninformed about their relative abilities. However, Proposition 3 and Corollary 3 suggest that with informed agents, organizational selection becomes more biased when merit depends not only on ability but also on efforts. Perhaps surprisingly, when viewed from this angle, our theory thus predicts a greater relevance of luck for selection in situations where agents carry a greater “responsibility” for their performance.

5 Societal luck

Our analysis in Sections 3 and 4 highlights the relevance of early career luck for an individual’s long-term success and explains how it is made persistent by organizational learning, even in environments where learning is very difficult. The “luck” on which we have focused so far derives from the inherent noisiness of individual performance. We have abstracted from factors that could impact agents systematically, such as the luck of possessing the “right identity” in the form of gender, race, ethnic origin, or socioeconomic background. There exists evidence showing that individuals with certain identities obtain advantages early in their career that can have long-lasting effects on social and economic outcomes.²¹ Conceptually, these forms of “societal luck” are different from what we have considered so

²¹Ciocca Eller (2023) provides evidence that differences in educational achievement of students attending colleges of *equal* selectivity can be traced to heterogeneous socioeconomic backgrounds. Bukodi et al. (2024) document the impact of “parental class” on attainment of ultra-elite scientific status in the UK.

far in that individuals might condition their actions on them. More specifically, agents' incentives might vary with their identity, and organizations might reward good performance with biases that depend on whether success was achieved with or without an exogenous advantage.

In this section, we extend our model of organizational learning by assuming that one agent $i \in \{A, B\}$ obtains an additive advantage of size $\alpha > 0$ that augments his initial performance $x_{i,1}$. This advantage is uncorrelated with ability and affects only the first-stage performance. In line with our earlier analysis, we examine to what extent the strategic interaction between the organization and the agents results in the societal luck of receiving a transitory advantage having a persistent effect. We define the *persistence of societal luck* as the probability that the initially advantaged agent is ultimately selected. For simplicity, our remaining analysis focuses on the case $q^0 = \frac{1}{2}$.²² Without loss of generality, we let agent A receive the advantage α , and we assume that this information is common knowledge. We distinguish two scenarios. Under *identity-dependent* (ID) biases, the principal can condition her bias β_i on the agents' identity $i \in \{A, B\}$. In contrast, under *identity-independent* (II) bias, the principal is required to set $\beta_A = \beta_B$. II bias might be a consequence of legislation aimed at preventing discriminatory practices.²³ Alternatively, there may be behavioral reasons why advantages are not accounted for, even when they are known to exist.²⁴ Our focus will be on how the persistence of societal luck differs between the equilibria in the regimes of ID and II biases.

In each scenario (ID or II), agents choose efforts optimally in response to the bias(es) they anticipate. For the by-now familiar reasons, second-stage efforts are identical across agents, so we can focus on the agents' first-stage effort differential $\Delta e_1 = e_{A,1} - e_{B,1}$. Suppose that under ID biases, agents' optimization results in $\Delta e_1 = \Delta e_1^*(\beta_A, \beta_B)$, and let $\Delta e_1 = \Delta e_1^*(\beta)$ be the analogous notation under II bias. As will become clear below, the principal's choice of bias depends on the anticipated *net advantage*, $\tilde{\alpha} = \alpha + \Delta e_1$, of the advantaged agent. Denote the principal's optimal choice of biases in the ID regime by $\beta_A^*(\tilde{\alpha})$ and $\beta_B^*(\tilde{\alpha})$, and let $\beta^*(\tilde{\alpha})$ be her optimal II bias. In the ID regime, an equilibrium is a combination of biases and net advantage $(\beta_A^{ID}, \beta_B^{ID}, \tilde{\alpha}^{ID})$ that are mutual best responses, that is, $\beta_A^{ID} = \beta_A^*(\tilde{\alpha}^{ID})$, $\beta_B^{ID} = \beta_B^*(\tilde{\alpha}^{ID})$, and $\tilde{\alpha}^{ID} = \alpha + \Delta e_1^*(\beta_A^{ID}, \beta_B^{ID})$. Similarly, in the

²²Our result in Corollary 4 on the persistence of societal luck generalizes to arbitrary q^0 , for small values of the exogenous advantage.

²³Title VII of the 1964 Civil Rights Act declares as "an unlawful employment practice [...] to discriminate against any individual because of his race, color, religion, sex, or national origin in admission to, or employment in, any program established to provide apprenticeship or other training."

²⁴Exley and Nielsen (2024) document that evaluators correctly expect women to be less confident than men in the assessment of their own abilities but fail themselves to account for this gender gap in their evaluations. We can show that, for small α , optimal II bias becomes insensitive to α , so that our analysis approximates the case where an advantage exists but is neglected by the principal.

II regime, an equilibrium $(\beta^{II}, \tilde{\alpha}^{II})$ satisfies $\beta^{II} = \beta^*(\tilde{\alpha}^{II})$ and $\tilde{\alpha}^{II} = \alpha + \Delta e_1^*(\beta^{II})$. We use Δe_1^{ID} and Δe_1^{II} to denote $\Delta e_1^*(\beta_A^{ID}, \beta_B^{ID})$ and $\Delta e_1^*(\beta^{II})$, respectively.

For arbitrary values of biases β_A and β_B , and net advantage $\tilde{\alpha}$, selective efficiency can be written as:

$$S(\beta_A, \beta_B, \tilde{\alpha}) = \frac{1}{2}[G(\lambda_1 h + \tilde{\alpha})G(\lambda_2 h + \beta_A) + G(-\lambda_1 h - \tilde{\alpha})G(\lambda_2 h - \beta_B)] \quad (13)$$

$$+ \frac{1}{2}[G(\lambda_1 h - \tilde{\alpha})G(\lambda_2 h + \beta_B) + G(-\lambda_1 h + \tilde{\alpha})G(\lambda_2 h - \beta_A)].$$

The terms in the first (respectively, second) square brackets are the probability that the better agent is selected conditional on being advantaged (respectively, disadvantaged). The principal's optimal ID biases β_A^* and β_B^* solve the first-order conditions

$$\frac{G(\lambda_1 h + \tilde{\alpha})}{G(-\lambda_1 h + \tilde{\alpha})} = \frac{g(\lambda_2 h - \beta_A^*)}{g(\lambda_2 h + \beta_A^*)} \quad \text{and} \quad \frac{G(\lambda_1 h - \tilde{\alpha})}{G(-\lambda_1 h - \tilde{\alpha})} = \frac{g(\lambda_2 h - \beta_B^*)}{g(\lambda_2 h + \beta_B^*)}. \quad (14)$$

The principal's optimal II bias β^* solves $\frac{\partial S}{\partial \beta} = 0$ under the constraint that $\beta_A = \beta_B = \beta$:

$$\frac{G(\lambda_1 h + \tilde{\alpha}) + G(\lambda_1 h - \tilde{\alpha})}{G(-\lambda_1 h + \tilde{\alpha}) + G(-\lambda_1 h - \tilde{\alpha})} = \frac{g(\lambda_2 h - \beta^*)}{g(\lambda_2 h + \beta^*)}. \quad (15)$$

Comparing (14) with (15) shows that in the II regime, the principal is restricted to set bias to match the ‘‘average’’ informativeness of a first-stage win, whereas ID biases allow the principal to adapt to whether such a win was achieved with or without a net advantage. From the log-concavity of g , it thus follows that for all $\tilde{\alpha}$:

$$\beta_B^*(\tilde{\alpha}) > \beta^*(\tilde{\alpha}) > \beta_A^*(\tilde{\alpha}) > 0. \quad (16)$$

Intuitively, a first-stage win against a net disadvantage is a stronger positive signal about the winner's ability than a first-stage win with a net advantage in the winner's favor. While comparing biases for *given* net advantage $\tilde{\alpha}$ is straightforward, a full comparison of the two scenarios requires a characterization of the agents' equilibrium effort differentials:

Proposition 4 (Incentive effects with societal luck) *Let $q^0 = \frac{1}{2}$ and suppose agent A's identity augments his first-stage performance by $\alpha > 0$.*

- (i) *Independently of whether bias can condition on agents' identity, in equilibrium agent A exerts a lower first-stage effort than agent B but maintains a strict net advantage:*

$$-\alpha < \Delta e_1^{II} < 0 \quad \text{and} \quad -\alpha < \Delta e_1^{ID} < 0.$$

(ii) If effort costs are $C_t(e_{i,t}) = \frac{c_t}{2}e_{i,t}^2$, with $c_t > 0$ sufficiently large for equilibrium in both regimes to be unique, and if the agents' ability difference h is sufficiently small, then

$$\tilde{\alpha}^{ID} = \alpha + \Delta e_1^{ID} < \alpha + \Delta e_1^{II} = \tilde{\alpha}^{II},$$

i.e. making bias identity-dependent reduces agent A's equilibrium net advantage.

Similarly to Section 4, where the first-stage competition was asymmetric due to agents' information about their relative abilities ($q^0 > \frac{1}{2}$), the agents' first-stage effort differential arises exclusively from the impact of the first-stage outcome on second-stage effort costs. Because a net advantage $\alpha + \Delta e_1 > 0$ makes A more likely to win the first stage, both the agents and the principal are less confident in the first-stage winner's ability when it is A compared to when it is B . Under II bias, the agents will therefore expect the biased second-stage competition to be more balanced, and consequently more costly, after a first-stage win by A than after a first-stage win by B . This difference in second-stage effort costs gives the advantaged agent, A , a weaker incentive than his rival to exert first-stage effort, resulting in $\Delta e_1 < 0$. Identity-dependent biases augment this "future effort-cost effect", reducing the induced Δe_1 further below zero, because the principal will optimally choose $\beta_A < \beta_B$ for any anticipated $\alpha + \Delta e_1 > 0$. As long as the ability difference h is not too large, a reduction in β_A makes the second-stage competition even more balanced following a win by A , and an increase in β_B makes the second-stage competition even less balanced following a win by B .²⁵

Having characterized the principal's and the agents' behavior under both II and ID biases, we are now ready to compare outcomes across these two regimes. Note first that, because the principal cannot commit to the level of bias, it is unclear a priori whether she will do better with ID biases, even though she is less constrained in this regime than under II bias. Yet it follows from Proposition 4(ii) that selective efficiency is always higher under ID biases than under II bias. This is because, by the envelope theorem, maximized selective efficiency under II bias is decreasing in the net advantage, and because under ID biases the net advantage is reduced. Thus, agents' effort responses *augment* the direct benefits of ID biases for selective efficiency.

Our main interest, however, is in the comparison of the persistence of societal luck,

²⁵The assumption of quadratic costs allows us to compare the agents' first-stage effort differential across the two regimes by focusing on the size of the difference in their marginal benefits of effort.

given by the probability P_α that the initially advantaged agent is ultimately selected:

$$P_\alpha(\beta_A, \beta_B, \tilde{\alpha}) = \frac{1}{2} \sum_{\Delta a \in \{-h, h\}} [G(\lambda_1 \Delta a + \tilde{\alpha})G(\lambda_2 \Delta a + \beta_A) + G(-\lambda_1 \Delta a - \tilde{\alpha})G(\lambda_2 \Delta a - \beta_B)]. \quad (17)$$

Under II bias, societal luck is *always* made persistent, that is, $P_\alpha(\beta^{II}, \beta^{II}, \alpha + \Delta e_1^{II}) > \frac{1}{2}$. This is most easily seen by noting that, for $\beta_A = \beta_B = \beta$, persistence can be written more simply as

$$P_\alpha(\beta, \beta, \tilde{\alpha}) = \frac{1}{2} \{1 + [G(\lambda_1 h + \tilde{\alpha}) - G(\lambda_1 h - \tilde{\alpha})][G(\lambda_2 h + \beta) - G(\lambda_2 h - \beta)]\}, \quad (18)$$

and in equilibrium, both the net advantage, $\tilde{\alpha}^{II} = \alpha + \Delta e_1^{II}$, and the principal's choice of bias, β^{II} , are strictly positive, as shown by Proposition 4 and (16). Intuitively, the advantaged agent is selected with higher probability than his rival, because he is more likely to win the first stage (despite his lower effort), and with II bias, the second stage is biased by the *same* amount, no matter the identity of the first-stage winner.

In striking contrast, allowing for ID biases may completely eliminate the persistence of societal luck. For example, we can show that, when the difference in agents' noise terms has a logistic distribution, then $P_\alpha(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) = \frac{1}{2}$ for *all* values of exogenous advantage $\alpha > 0$.²⁶ The following corollary to Proposition 4 provides a general comparison of the persistence of societal luck between the equilibria in the II and ID regimes. It also compares the expected utility difference ΔU between the advantaged and the disadvantaged agent across these equilibria:

$$\Delta U(\beta_A, \beta_B, \alpha + \Delta e_1) \equiv [2P_\alpha(\beta_A, \beta_B, \alpha + \Delta e_1) - 1] - [C_1(e_{A,1}) - C_1(e_{B,1})]. \quad (19)$$

Corollary 4 (Persistence of societal luck) *Under the assumptions of Proposition 4 (ii) and for exogenous advantage $\alpha > 0$ sufficiently small, allowing bias to be identity-dependent*

(i) *reduces the persistence of societal luck:*

$$P_\alpha(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) < P_\alpha(\beta^{II}, \beta^{II}, \alpha + \Delta e_1^{II})$$

²⁶The logistic distribution does not satisfy part (iv) of Assumption 1. Yet, all our results go through as long as the first stage is not too informative relative to the second one, that is, λ_1 is not too high relative to λ_2 .

(ii) reduces the expected utility difference between agents:

$$\Delta U(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) < \Delta U(\beta_A^{II}, \beta_B^{II}, \alpha + \Delta e_1^{II}).$$

We stress that ID biases reduce the persistence of societal luck via two distinct channels. For any *given* net advantage $\tilde{\alpha}$, persistence is reduced by ID biases because $\beta_A^*(\tilde{\alpha}) < \beta^*(\tilde{\alpha}) < \beta_B^*(\tilde{\alpha})$, and P_α in (17) is increasing in β_A and decreasing in β_B . This reduction in persistence reflects the fact that, whatever the first-stage outcome, ID biases effectively penalize in the second stage the agent who benefited from the first-stage advantage.

The second driver behind the reduced persistence is the effect ID biases have on agents' incentives. As shown by Proposition 4(ii), ID biases induce the disadvantaged agent to compensate *even more* for his disadvantage through higher first-stage effort than under II bias. This generates a further reduction in the persistence of societal luck, as long as α is sufficiently small, because for small $\tilde{\alpha}$, $P_\alpha(\beta^*(\tilde{\alpha}), \beta^*(\tilde{\alpha}), \tilde{\alpha})$ is increasing in $\tilde{\alpha}$.²⁷

In summary, the results in this section can be interpreted as offering support for *affirmative action*, in the form of selection processes that allow organizations to condition biases on agents' identities—for instance, through gender-specific mentoring or grants accounting for socioeconomic backgrounds. Roemer (2000) argues that an *equal opportunity principle* should be applied at the entry level of careers, e.g. for admissions to medical school, while a *non-discrimination principle* should govern the selection for final positions, such as the licensing of surgeons. The first part of Corollary 4 shows that applying a non-discrimination principle (requiring bias to be identity-independent) in the selection for later positions can backfire, by propagating disadvantages stemming from a failure to establish equal opportunity upon entry. Moreover, the second part of Corollary 4 shows that in our setting, affirmative action is in fact doubly beneficial, because identity-dependent biases also decrease inequality. Finally, our analysis highlights that the incentive effects of such policies do not hinder their effectiveness but rather amplify their benefits.²⁸

²⁷Under cardinal information, ID biases would allow the principal to completely filter out the effect of α . As a result, the advantaged agent would be no more likely to be selected than his rival, and the unique equilibrium first-stage effort differential would be zero.

²⁸Analysis of the case where $q^0 > \frac{1}{2}$ indicates that the beneficial effects of agents' effort responses are further amplified when agents are informed about their relative abilities.

6 Conclusion

When the careers of National Hockey League players (Deaner et al., 2013) or S&P 500 CEOs (Du et al., 2012) are kick-started by the proximity of their birthdays to a cut-off or when hedge funds (Cong and Xiao, 2022) or venture capitalists (Nanda et al., 2020) persistently outperform the market following a fortunate initial investment, luck seems to play an unjustified role in the selection of the most gifted. Such findings and related anecdotes might seem to support recent critiques of a meritocratic worldview (e.g. Piketty, 2014; Sandel, 2020). Their argument is that, in spite of this worldview forming the basis of modern democratic societies, meritocracy is a myth, used to justify exorbitant degrees of economic and social inequality. The main contribution of this paper is to show that making initial luck have a persistent effect on selection is consistent with—if not a necessary feature of—a society aiming to allocate resources and decision power to the most able individuals.

Our theory illuminates a basic mechanism behind inequality by rationalizing the persistence of luck as an equilibrium outcome of the strategic interaction between an organization aiming to maximize selective efficiency and heterogeneous agents capable of influencing their likelihood of being selected through costly efforts. We have characterized the settings where the impact of initial luck can be expected to be most amplified. This happens when agents are informed about their relative abilities and the organization is restricted to use ordinal rather than cardinal performance information. Both conditions seem more likely to be met towards the top of an organization’s hierarchy, which suggests that luck is a more significant determinant of selection where, arguably, selection matters most, both for the organization and for the induced inequality among the agents.

We have also analyzed how organizational learning impacts the persistence of a different type of luck—termed “societal luck”—which reflects advantages that individuals derive from their identities. Our results suggest that non-discrimination policies that constrain an organization’s selection process may backfire, by propagating disadvantages stemming from unequal initial opportunities, especially when agents strategically respond to such policies.

Appendix

Proof of Lemmas 1 and 2

Use superscripts w and l , respectively, to distinguish the cases where agent A won and lost the first stage. Define $\Delta e_1 = e_{A,1} - e_{B,1}$, $\Delta e_2^w = e_{A,2}^w - e_{B,2}^w$, and $\Delta e_2^l = e_{A,2}^l - e_{B,2}^l$. Let $q^w(\Delta e_1, q^0)$ and $q^l(\Delta e_1, q^0)$ denote the posterior probabilities that the *winner* of the first stage is the more able agent, given q^0 and Δe_1 . When there is no risk of confusion, we suppress the arguments of the posteriors.

We first show that agents exert identical effort in the second stage and that this holds *independently* of q^0 and Δe_1 . In case w , A 's and B 's first-order conditions determining second-stage efforts are:

$$\begin{aligned} C_2'(e_{A,2}^w) &= q^w g(h + \beta + \Delta e_2^w) + (1 - q^w)g(-h + \beta + \Delta e_2^w) \\ C_2'(e_{B,2}^w) &= q^w g(-h - \beta - \Delta e_2^w) + (1 - q^w)g(h - \beta - \Delta e_2^w). \end{aligned}$$

By the symmetry of g , the marginal returns to effort are identical, so $e_{A,2}^w = e_{B,2}^w$. An analogous argument for case l shows that $e_{A,2}^l = e_{B,2}^l$.

Now consider the agents' incentives for first-stage efforts. We can write the overall utility of agent A as follows:

$$\begin{aligned} -C_1(e_{A,1}) &+ q^0 \{G(\lambda_1 h + \Delta e_1) [G(\lambda_2 h + \beta + \Delta e_2^w) - C_2(e_{A,2}^w)] \\ &+ [1 - G(\lambda_1 h + \Delta e_1)] [G(\lambda_2 h - \beta + \Delta e_2^l) - C_2(e_{A,2}^l)]\} \\ &+ (1 - q^0) \{G(-\lambda_1 h + \Delta e_1) [G(-\lambda_2 h + \beta + \Delta e_2^w) - C_2(e_{A,2}^w)] \\ &+ [1 - G(-\lambda_1 h + \Delta e_1)] [G(-\lambda_2 h - \beta + \Delta e_2^l) - C_2(e_{A,2}^l)]\}. \end{aligned}$$

A change in $e_{A,1}$ does not affect $e_{B,2}^w$, $e_{B,2}^l$, or β , because it is unobservable, and the local effect via the induced changes in $e_{A,2}^w$ and $e_{A,2}^l$ is zero by the envelope theorem. Using $\Delta e_2^w = \Delta e_2^l = 0$ and the symmetry of g around 0, the first-order condition for $e_{A,1}$ can be written as

$$\begin{aligned} C_1'(e_{A,1}) &= [q^0 g(\lambda_1 h + \Delta e_1) + (1 - q^0)g(-\lambda_1 h + \Delta e_1)] \\ &\cdot \{[G(\lambda_2 h + \beta) - G(\lambda_2 h - \beta)] - [C_2(e_2^w) - C_2(e_2^l)]\} \end{aligned} \tag{20}$$

Analogously, for agent B the first-order condition for $e_{B,1}$ can be written as

$$\begin{aligned} C_1'(e_{B,1}) &= [(1 - q^0)g(\lambda_1 h - \Delta e_1) + q^0 g(-\lambda_1 h - \Delta e_1)] \\ &\cdot \{[G(\lambda_2 h + \beta) - G(\lambda_2 h - \beta)] + [C_2(e_2^w) - C_2(e_2^l)]\} \end{aligned} \tag{21}$$

Again using the symmetry of g , and noting that the component of the marginal benefit of the first-stage effort stemming from the enhanced probability of selection is identical for the two agents, when we subtract (21) from (20), we get

$$\frac{C'_1(e_{A,1}) - C'_1(e_{B,1})}{C_2(e_2^l) - C_2(e_2^w)} = 2 [q^0 g(\lambda_1 h + \Delta e_1) + (1 - q^0) g(-\lambda_1 h + \Delta e_1)] \quad (22)$$

Given that costs are strictly increasing and strictly convex, we conclude that in equilibrium, $\Delta e_1 = e_{A,1} - e_{B,1}$ and $e_2^l - e_2^w$ must have the same sign.

To determine the sign of $e_2^l - e_2^w$, compare agent A 's first-order conditions for the second-stage effort, after a first-stage win by A vs. after a first-stage loss by A , respectively:

$$C'_2(e_2^w) = q^w(\Delta e_1, q^0) g(\lambda_2 h + \beta) + (1 - q^w(\Delta e_1, q^0)) g(-\lambda_2 h + \beta), \quad (23)$$

$$C'_2(e_2^l) = q^l(\Delta e_1, q^0) g(-\lambda_2 h - \beta) + (1 - q^l(\Delta e_1, q^0)) g(\lambda_2 h - \beta). \quad (24)$$

Subtracting the second FOC from the first, and using the symmetry of g , gives

$$C'_2(e_2^w) - C'_2(e_2^l) = [q^w(\Delta e_1, q^0) - q^l(\Delta e_1, q^0)] [g(\lambda_2 h + \beta) - g(-\lambda_2 h + \beta)]. \quad (25)$$

The strict log-concavity and symmetry of g imply that for any $\beta > 0$, $g(\lambda_2 h + \beta) - g(-\lambda_2 h + \beta) < 0$, while for $\beta = 0$, $g(\lambda_2 h + \beta) - g(-\lambda_2 h + \beta) = 0$. Hence, since costs are strictly convex,

$$e_2^l - e_2^w \geq 0 \iff q^w(\Delta e_1, q^0) - q^l(\Delta e_1, q^0) \geq 0. \quad (26)$$

The posterior beliefs $q^w(\Delta e_1, q^0)$ and $q^l(\Delta e_1, q^0)$ are given by

$$q^w(\Delta e_1, q^0) = \frac{q^0 G(\lambda_1 h + \Delta e_1)}{q^0 G(\lambda_1 h + \Delta e_1) + (1 - q^0) G(-\lambda_1 h + \Delta e_1)}, \quad (27)$$

$$q^l(\Delta e_1, q^0) = \frac{(1 - q^0) G(\lambda_1 h - \Delta e_1)}{(1 - q^0) G(\lambda_1 h - \Delta e_1) + q^0 G(-\lambda_1 h - \Delta e_1)} \quad (28)$$

Observe that q^w and q^l are, respectively, strictly decreasing and strictly increasing in Δe_1 . For $q^0 = \frac{1}{2}$, they are equal at $\Delta e_1 = 0$, while for $q^0 > \frac{1}{2}$, they are equal at some $\Delta e_1 > 0$.

We are now in a position to complete the proof of Lemma 1. Let $q^0 = \frac{1}{2}$. Suppose first that agents anticipate bias $\beta \leq 0$. Then, the agents would like to decrease their first-stage performance (if $\beta < 0$) or are indifferent with respect to it (if $\beta = 0$) while the effort is costly. Hence, $e_{A,1} = e_{B,1} = 0$. Now, let agents anticipate bias $\beta > 0$. Suppose, for contradiction, that $\Delta e_1 > 0$. Then $q^w(\Delta e_1, q^0) - q^l(\Delta e_1, q^0) < 0$, so by (26), $e_2^l < e_2^w$.

In turn, this implies, using (22), that $\Delta e_1 < 0$, which is a contradiction. Analogously, assuming that $\Delta e_1 < 0$ would also lead to a contradiction. Hence, equilibrium requires equal first-stage efforts: $e_{A,1} = e_{B,1}$. These are unique since with $\Delta e_1 = 0$, the right-hand sides of (20) and (21) are independent of the common level of e_1 .

To complete the proof of Lemma 2, we need to show that for any $q^0 > \frac{1}{2}$ and any $\beta > 0$, equilibrium entails $e_{A,1} - e_{B,1} > 0$. Suppose, for contradiction, that $\Delta e_1 \leq 0$. Then from (27) and (28), $q^w(\Delta e_1, q^0) - q^l(\Delta e_1, q^0) > 0$, because the agents' prior is that A is more able *and* a first-stage win by A despite an effort disadvantage is per se a stronger signal of ability than a first-stage win by B with an effort advantage. By (26), it follows that $e_2^l > e_2^w$. In turn, this implies, using (22), that $\Delta e_1 > 0$, which is a contradiction. ■

Proof of Proposition 1

Equilibrium bias maximizes selective efficiency, $S(\beta; h)$, which for $q_0 = \frac{1}{2}$ by Lemma 1 is given by (1). We use sub-indices to denote partial derivatives. For any $h > 0$, Assumption 1 ensures that the first-order condition $S_\beta(\beta; h) = 0$ uniquely determines the optimal bias $\beta^*(h)$:

$$S_\beta(\beta^*(h); h) = G(\lambda_1 h) g(\lambda_2 h + \beta^*(h)) - [1 - G(\lambda_1 h)] g(\lambda_2 h - \beta^*(h)) = 0.$$

To see that $\beta^*(h) > 0$ for all $h > 0$ note that $G(\lambda_1 h) > 1 - G(\lambda_1 h)$. However, $\lim_{h \rightarrow 0} S_\beta(\beta, h) = 0 \forall \beta$. Characterizing $\beta_0^* \equiv \lim_{h \rightarrow 0} \beta^*(h)$ thus requires totally differentiating $S_\beta(\beta^*(h); h)$ with respect to h , setting it equal to 0, and letting $h \rightarrow 0$. Total differentiation yields

$$\frac{d}{dh} S_\beta(\beta^*(h); h) = S_{\beta h}(\beta^*(h); h) + S_{\beta\beta}(\beta^*(h); h) \frac{\partial \beta^*(h)}{\partial h}, \quad (29)$$

where $\lim_{h \rightarrow 0} S_{\beta\beta}(\beta; h) = 0 \forall \beta$ (since $\lim_{h \rightarrow 0} S_\beta(\beta; h) = 0 \forall \beta$). Hence, (29) and Assumption 1(i) imply that β_0^* solves

$$\lim_{h \rightarrow 0} S_{\beta h}(\beta^*(h); h) = S_{\beta h}(\beta_0^*, 0) = 2\lambda_1 g(0)g(\beta_0^*) + \lambda_2 g'(\beta_0^*) = 0,$$

which gives (3). Since Assumptions 1(i) and 1(iii) guarantee that $L(0) = 0$ and that L is strictly increasing, it follows that $\beta_0^* > 0$. ■

Proof of Proposition 2

To abbreviate notation we let $k = |\Delta x_1| \geq 0$ denote the observed first-stage margin of victory.

Part (i) Having observed the margin of victory, k , the principal chooses β to maximize the objective in (6), and the first-order condition is

$$S_{\beta}^{card}(\beta, k; h) = g(k - \lambda_1 h)g(\lambda_2 h + \beta) - g(k + \lambda_1 h)g(\lambda_2 h - \beta) = 0. \quad (30)$$

By Assumption 1, (30) uniquely determines the optimal cardinal bias $\beta^{card}(k, h)$ as a strictly increasing function of k , equal to zero for $k = 0$. Since $\lim_{h \rightarrow 0} S_{\beta}^{card}(\beta, k; h) = 0 \forall \beta, k$, characterizing $\beta_0^{card}(k) \equiv \lim_{h \rightarrow 0} \beta^{card}(k, h)$ requires totally differentiating the value $S_{\beta}^{card}(\beta^{card}(k, h), k; h)$ with respect to h , setting it equal to zero, and letting $h \rightarrow 0$. Doing so shows that $\beta_0^{card}(k)$ solves $\lim_{h \rightarrow 0} S_{\beta h}^{card}(\beta, k; h) = 0$, which yields

$$L(\beta_0^{card}(k)) = \frac{\lambda_1}{\lambda_2} L(k),$$

which is equation (7). By Assumption 1, $L(0) = 0$ and $L(k) > 0 \forall k > 0$. Hence, $\beta_0^{card}(k) > 0 \forall k > 0$.

Part (ii) Given (3) and (7), we need only show that $\mathbb{E}[L(k)] = 2g(0)$. As $h \rightarrow 0$, the density of k converges to $2g(k)$ on support $[0, z]$. Hence

$$\mathbb{E}[L(k)] = \int_0^z L(k)2g(k)dk = -2 \int_0^z g'(k)dk = 2g(0),$$

using $g(z) = 0$, which is implied by Assumption 1(iii). ■

Proof of Corollary 2

Part (i) When L is convex, (8) implies that $\beta_0^* \geq E[\beta_0^{card}(k)]$ and hence

$$G(\beta_0^*) \geq G(E[\beta_0^{card}(k)]), \quad (31)$$

since G is strictly increasing. Strict log-concavity and symmetry of g imply that G is strictly concave on the positive domain, so

$$G(E[\beta_0^{card}(k)]) > E[G(\beta_0^{card}(k))]. \quad (32)$$

Inequalities (31) and (32) together imply (10).

Part (ii) Whichever type of information, ordinal or cardinal, is used, and given the ex ante symmetry of the selection process with respect to agents A and B , the limiting

value of persistence as $h \rightarrow 0$ can be expressed as

$$\begin{aligned}
& 2\mathbb{P}(\text{select } A, \Delta\epsilon_1 > 0) \\
&= 2 [\mathbb{P}(\text{select } A, \Delta\epsilon_1 > 0, \Delta\epsilon_2 > 0) + \mathbb{P}(\text{select } A, \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0)] \\
&= \frac{1}{2} [\mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 > 0) + \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0)],
\end{aligned}$$

where we have used the fact that $\mathbb{P}(\Delta\epsilon_1 > 0, \Delta\epsilon_2 > 0) = \mathbb{P}(\Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0) = \frac{1}{4}$. Since

$$\mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 > 0, \text{ord.}) = \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 > 0, \text{card.}) = 1, \quad (33)$$

it follows that $P_0^* > P_0^{\text{card}}$ if and only if

$$\mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{ord.}) > \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{card.}). \quad (34)$$

Whether ordinal or cardinal information is used, the ex ante symmetry of the selection process with respect to A and B means that the ex ante probability of selecting A is $\frac{1}{2}$. Using the first equality in (33), and the fact that

$$\mathbb{P}(\text{select } A \mid \Delta\epsilon_1 < 0, \Delta\epsilon_2 < 0, \text{ord.}) = \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 < 0, \Delta\epsilon_2 < 0, \text{card.}) = 0, \quad (35)$$

it thus must be that

$$\begin{aligned}
& \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{ord.}) + \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 < 0, \Delta\epsilon_2 > 0, \text{ord.}) \\
&= \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 > 0, \Delta\epsilon_2 < 0, \text{card.}) + \mathbb{P}(\text{select } A \mid \Delta\epsilon_1 < 0, \Delta\epsilon_2 > 0, \text{card.}).
\end{aligned} \quad (36)$$

Using (36), it is then straightforward to confirm that (34) holds if and only if (11) is satisfied. ■

Proof of Proposition 3

In the limit as $h \rightarrow 0$, (25) implies that $e_2^w - e_2^l \rightarrow 0$, since agents' posterior beliefs about their relative ability become irrelevant to their second-stage effort incentives. The first-stage effort differential $\Delta e_1 = e_{A,1} - e_{B,1}$ therefore approaches 0 as $h \rightarrow 0$, since the right-hand sides of (20) and (21) become equal.

Using this result, we now characterize the principal's optimal choice of bias, for any anticipated first-stage effort differential Δe_1 . The principal chooses β to maximize selective

efficiency $S(\beta, h, q^0)$, where

$$\begin{aligned} S(\beta; h, q^0) &= [q^0 G(\lambda_1 h + \Delta e_1) + (1 - q^0) G(\lambda_1 h - \Delta e_1)] G(\lambda_2 h + \beta) \\ &+ [q^0 G(-\lambda_1 h - \Delta e_1) + (1 - q^0) G(-\lambda_1 h + \Delta e_1)] G(\lambda_2 h - \beta). \end{aligned} \quad (37)$$

The first-order condition for β is

$$\begin{aligned} S_\beta(\beta; h, q^0) &= [q^0 G(\lambda_1 h + \Delta e_1) + (1 - q^0) G(\lambda_1 h - \Delta e_1)] g(\lambda_2 h + \beta) \\ &- [q^0 (1 - G(\lambda_1 h + \Delta e_1)) + (1 - q^0) (1 - G(\lambda_1 h - \Delta e_1))] g(\lambda_2 h - \beta) = 0. \end{aligned} \quad (38)$$

Since $\lim_{h \rightarrow 0} \Delta e_1 = 0$, $\lim_{h \rightarrow 0} S_\beta(\beta; h, q^0) = 0$ for all β . As in the proof of Proposition 1, characterizing the optimal bias $\beta^*(h)$ in the limit as $h \rightarrow 0$ thus requires totally differentiating $S_\beta(\beta^*(h), h, q^0)$ with respect to h , setting it equal to 0, and letting $h \rightarrow 0$. Since $\lim_{h \rightarrow 0} S_\beta(\beta; h, q^0) = 0$ for all β , $\lim_{h \rightarrow 0} S_{\beta\beta}(\beta; h, q^0)$ for all β . Hence the limiting optimal bias as $h \rightarrow 0$, β_0^* , solves the first-order condition

$$\begin{aligned} 0 = \lim_{h \rightarrow 0} S_{\beta h}(\beta^*(h); h, q^0) &= S_{\beta h}(\beta_0^*; 0, q^0) \\ &= 2g(0)g(\beta_0^*) \left[\lambda_1 + (2q^0 - 1) \frac{\partial \Delta e_1}{\partial h} \Big|_{h \rightarrow 0} \right] + \lambda_2 g'(\beta_0^*). \end{aligned} \quad (39)$$

To complete the characterization of equilibrium in the limit as $h \rightarrow 0$, we must determine how the limiting derivative with respect to h of the agents' best-response effort differential, $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$, depends on their anticipations about the principal's choice of β . The derivation of $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$ is simplified by the following observation, which is based on a symmetry argument: $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$.

To show that $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$, we begin by observing that since $\lim_{h \rightarrow 0} \beta^*(h)$ solves the first-order condition $\lim_{h \rightarrow 0} S_{\beta h} = 0$, the sign of $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$ will be determined by the sign of $\lim_{h \rightarrow 0} S_{\beta h h}(\beta^*(h); h, q^0)$. We will show that $\lim_{h \rightarrow 0} S_{hh}(\beta; h, q^0) = 0$ for all β, q^0 , from which it follows that $\lim_{h \rightarrow 0} S_{\beta h h}(\beta^*(h); h, q^0) = 0$ for all β, q^0 and therefore $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$.

To prove that $\lim_{h \rightarrow 0} S_{hh}(\beta; h, q^0) = 0$ for all β, q^0 , we will show that, for any $\beta, S(\beta; h, q^0)$, regarded as a function of $h \in \mathfrak{R}$, displays 180° rotational symmetry around the point $(h = 0, S = \frac{1}{2})$, that is, $S(\beta; h, q^0) = 1 - S(\beta; -h, q^0)$. To interpret the mathematical expression $S(\beta; -h, q^0)$, temporarily set $\Delta e_1 = 0$; $S(\beta; -h, q^0)$ then gives the probability of selecting the more able agent when the principal assigns bias β in favor of the first-stage *loser*. In such a setting, the endogenous first-stage effort differential would switch sign, that is, $\Delta e_1(-h) = -\Delta e_1(h)$, as can be seen from (22) and (25). Using $\Delta e_1(-h) = -\Delta e_1(h)$,

we have

$$\begin{aligned} S(\beta; -h, q^0) &= [q^0 G(\lambda_1 h + \Delta e_1(h)) + (1 - q^0) G(\lambda_1 h - \Delta e_1(h))] G(-\lambda_2 h - \beta) \\ &+ [q^0 G(-\lambda_1 h - \Delta e_1(h)) + (1 - q^0) G(-\lambda_1 h + \Delta e_1(h))] G(-\lambda_2 h + \beta). \end{aligned} \quad (40)$$

It follows from (40) and (37) that for all h, β, q^0 , $S(\beta; h, q^0) = 1 - S(\beta; -h, q^0)$. Differentiating this identity twice with respect to h and letting $h \rightarrow 0$ then yields $\lim_{h \rightarrow 0} S_{hh}(\beta; h, q^0) = 0$ for all β, q^0 .

Having established that $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$, we now return to the analysis of how $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$ depends on agents' anticipations about the principal's choice of β . Differentiating the agents' first-order conditions for first-stage effort, (20) and (21), with respect to h , letting $h \rightarrow 0$, and using $\lim_{h \rightarrow 0} \frac{\partial \beta^*(h)}{\partial h} = 0$, yields

$$C_1'''(e_0) \left[\frac{\partial e_{A,1}}{\partial h} - \frac{\partial e_{B,1}}{\partial h} \right] = -4\lambda_2 g(0) X'(g(\beta)) g'(\beta) (2q^0 - 1), \quad (41)$$

where e_0 is the agents' common limiting first-stage effort, given β , which solves $C_1'(e_0) = g(0)[2G(\beta) - 1]$, and the function $X(\cdot) \equiv C_2(C_2'^{-1}(\cdot))$. Note that e_0 is independent of q^0 and that strict convexity of $C_2(\cdot)$ ensures that $X(\cdot)$ is strictly increasing. For any anticipated $\beta > 0$ and any $q^0 > \frac{1}{2}$, the right-hand side of (41) is strictly positive, so $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h} > 0$.

An *equilibrium* value of β as $h \rightarrow 0$, β_0^* , solves the fixed-point equation derived from (39), recognizing the dependence of $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h}$ on β_0 :

$$2g(0) \left[\lambda_1 + (2q^0 - 1) \frac{\partial \Delta e_1(\beta_0^*; 0, q^0)}{\partial h} \right] = \lambda_2 L(\beta_0^*). \quad (42)$$

Since $\lim_{h \rightarrow 0} \frac{\partial \Delta e_1}{\partial h} > 0$ for all $\beta > 0, q^0 > \frac{1}{2}$, the left-hand side of (42) is strictly positive, so any fixed point β_0^* must be strictly positive. To show that the fixed point is unique, use (41) to substitute for $\frac{\partial \Delta e_1(\beta_0; 0, q^0)}{\partial h}$ in (42). This yields, after rearrangement,

$$2\lambda_1 g(0) = \lambda_2 L(\beta_0^*) \left[1 - \frac{8(g(0))^2}{C_1'''(e_0)} X'(g(\beta_0^*)) g(\beta_0^*) (2q^0 - 1)^2 \right]. \quad (43)$$

For $C_t(e_{i,t}) = \frac{c_t}{2} e_{i,t}^2$, $C_1'''(e_0)$ is a constant, and $X'(\cdot)$ is linear, so the expression in square brackets on the right-hand side of (43) is strictly increasing in β_0 . For quadratic costs, therefore, the right-hand side of (43) is strictly increasing in β_0 whenever the expression in square brackets is positive. Since the left-hand side of (43) is strictly positive, there is a unique equilibrium value β_0^* . Finally, since the right-hand side of (43) is strictly decreasing

in q^0 for all $\beta_0 > 0$, the equilibrium β_0^* is increasing in q^0 . ■

Proof of Proposition 4

We first derive properties of the principal's optimal bias, given her belief (correct in equilibrium) about the agents' effort differential and the corresponding net advantage $\tilde{\alpha}$. First note that, given net advantage $\tilde{\alpha}$, the principal's optimal biases $\beta_A^*(\tilde{\alpha})$, $\beta_B^*(\tilde{\alpha})$, and $\beta^*(\tilde{\alpha})$ are strictly positive. This is because the left hand sides of the first-order conditions (14) and (15) are strictly larger than one, while the right hand sides are equal to one when bias is zero and strictly increasing in bias by the log-concavity of g . Moreover, $\beta_A^*(\tilde{\alpha}) < \beta^*(\tilde{\alpha}) < \beta_B^*(\tilde{\alpha})$ for $\tilde{\alpha} > 0$ because the principal's "confidence" in the first-stage winner's ability is strictly decreasing in his net advantage, i.e.

$$\frac{G(\lambda_1 h + \tilde{\alpha})}{G(-\lambda_1 h + \tilde{\alpha})} > \frac{G(\lambda_1 h + \tilde{\alpha}) + G(\lambda_1 h - \tilde{\alpha})}{G(-\lambda_1 h + \tilde{\alpha}) + G(-\lambda_1 h - \tilde{\alpha})} > \frac{G(\lambda_1 h - \tilde{\alpha})}{G(-\lambda_1 h - \tilde{\alpha})}. \quad (44)$$

For the same reason, $\beta_A^*(\tilde{\alpha})$ and $\beta^*(\tilde{\alpha})$ are strictly decreasing whereas $\beta_B^*(\tilde{\alpha})$ is strictly increasing. As all three terms in (44) converge to $\frac{G(\lambda_1 h)}{G(-\lambda_1 h)}$ for $\tilde{\alpha} \rightarrow 0$, it holds that $\lim_{\alpha \rightarrow 0} \beta_A^{ID} = \lim_{\alpha \rightarrow 0} \beta_B^{ID} = \lim_{\alpha \rightarrow 0} \beta^{II}$. Finally, differentiating the left hand side of (15) with respect to $\tilde{\alpha}$ gives

$$\frac{2[g(\lambda_1 h + \tilde{\alpha}) - g(\lambda_1 h - \tilde{\alpha})]}{[G(-\lambda_1 h + \tilde{\alpha}) + G(-\lambda_1 h - \tilde{\alpha})]^2},$$

which converges to zero for $\alpha \rightarrow 0$, proving that $\lim_{\tilde{\alpha} \rightarrow 0} \frac{\partial \beta^*}{\partial \tilde{\alpha}} = 0$. And since the first-order conditions (14) determining optimal identity-dependent biases are identical except for the sign of $\tilde{\alpha}$, it has to hold that $\lim_{\tilde{\alpha} \rightarrow 0} \frac{\partial \beta_A^*}{\partial \tilde{\alpha}} = -\lim_{\tilde{\alpha} \rightarrow 0} \frac{\partial \beta_B^*}{\partial \tilde{\alpha}}$. Having established the properties of the optimal bias response we can now turn our attention to the claims in Proposition 4 about agents' first-stage effort differential.

Part (i) The proof of this claim treats jointly the cases of identity-dependent and identity-independent bias, for the latter simply impose $\beta_A = \beta_B = \beta$ and all arguments go through for all $\beta > 0$. In the second stage, agent A 's effort equals agent B 's effort. The proof of this claim is analogous to that of Lemma 1 and will thus be omitted. Let e_2^w and e_2^l denote the agents' (identical) second-stage efforts after the advantaged agent A won or lost the first stage, respectively. Agent A 's expected utility in stage one is then given by

$$-C_1(e_{A,1}) + \frac{1}{2} \sum_{\Delta a \in \{-h, h\}} \{G(\lambda_1 \Delta a + \alpha + \Delta e_1)[G(\lambda_2 \Delta a + \beta_A) - C_2(e_2^w)] \\ + G(-\lambda_1 \Delta a - \alpha - \Delta e_1)[G(\lambda_2 \Delta a - \beta_B) - C_2(e_2^l)]\},$$

and the corresponding first-order condition is

$$2C'_1(e_{A,1}) = \sum_{\Delta a \in \{-h, h\}} g(\lambda_1 \Delta a + \alpha + \Delta e_1) [G(\lambda_2 \Delta a + \beta_A) - G(\lambda_2 \Delta a - \beta_B) + C_2(e_2^l) - C_2(e_2^w)]. \quad (45)$$

Similarly, for agent B we get

$$2C'_1(e_{B,1}) = \sum_{\Delta a \in \{-h, h\}} g(\lambda_1 \Delta a + \alpha + \Delta e_1) [G(-\lambda_2 \Delta a + \beta_B) - G(-\lambda_2 \Delta a - \beta_A) + C_2(e_2^w) - C_2(e_2^l)]. \quad (46)$$

Comparing the marginal benefits of effort across agents, it follows from $G(x) = 1 - G(x)$ that those parts stemming from the enhanced probability of selection are identical. Subtracting B 's first-order condition from A 's yields:

$$\frac{C'_1(e_{A,1}) - C'_1(e_{B,1})}{C_2(e_2^l) - C_2(e_2^w)} = \sum_{\Delta a \in \{-h, h\}} g(\lambda_1 \Delta a + \alpha + \Delta e_1). \quad (47)$$

Given that C_1 and C_2 are increasing and convex, in equilibrium $\Delta e_1 = e_{A,1} - e_{B,1}$ and $e_2^l - e_2^w$ must have the same sign. To determine the sign of $e_2^l - e_2^w$, consider the advantaged agent A 's expected utility in the second stage, separately for the two cases where the advantaged agent won (w) or lost (l) the first stage, respectively:

$$\begin{aligned} q^w G(\lambda_2 h + \beta_A + \Delta e_2^w) + (1 - q^w) G(-\lambda_2 h + \beta_A + \Delta e_2^w) - C_2(e_{A,2}^w), \\ q^l G(-\lambda_2 h - \beta_B + \Delta e_2^l) + (1 - q^l) G(\lambda_2 h - \beta_B + \Delta e_2^l) - C_2(e_{A,2}^l). \end{aligned}$$

Here we have introduced

$$q^w = \frac{G(\lambda_1 h + \alpha + \Delta e_1)}{G(\lambda_1 h + \alpha + \Delta e_1) + G(-\lambda_1 h + \alpha + \Delta e_1)}, \quad (48)$$

$$q^l = \frac{G(\lambda_1 h - \alpha - \Delta e_1)}{G(\lambda_1 h - \alpha - \Delta e_1) + G(-\lambda_1 h - \alpha - \Delta e_1)} \quad (49)$$

to denote the posterior probabilities that the winner of the first-stage is the more able agent. The corresponding first-order conditions determining e_2^w and e_2^l are

$$C'_2(e_2^w) = q^w g(\lambda_2 h + \beta_A) + (1 - q^w) g(-\lambda_2 h + \beta_A), \quad (50)$$

$$C'_2(e_2^l) = q^l g(-\lambda_2 h - \beta_B) + (1 - q^l) g(\lambda_2 h - \beta_B). \quad (51)$$

Note that q^w (resp. q^l) is a decreasing (resp. increasing) function of the net advantage

and that

$$q^l > q^w \Leftrightarrow \alpha + \Delta e_1 > 0.$$

Also note that, as we argued above, with identity-dependent biases, the principal awards a larger bias when she is more certain that the first-stage winner is the more able agent, that is, in equilibrium $\beta_A - \beta_B$ and $q^w - q^l$ have the same sign.

We now argue, by contradiction, that $-\alpha < \Delta e_1 < 0$. Suppose, instead, that, $\Delta e_1 \leq -\alpha$. Then $\alpha + \Delta e_1 \leq 0$ implies that $q^l \leq q^w$ and thus $\beta_A \geq \beta_B$. (For identity-independent bias, this condition holds trivially.) We have, for all $\beta \in (0, \beta_A]$,

$$\frac{q^w}{1 - q^w} = \frac{g(\lambda_2 h - \beta_A)}{g(\lambda_2 h + \beta_A)} \geq \frac{g(\lambda_2 h - \beta)}{g(\lambda_2 h + \beta)} > \frac{g'(\lambda_2 h - \beta)}{g'(\lambda_2 h + \beta)} = -\frac{g'(-\lambda_2 h + \beta)}{g'(\lambda_2 h + \beta)},$$

where the first equality is the principal's first-order condition for β_A , the two inequalities follow from $\beta \in (0, \beta_A]$ and the strict log-concavity of g , and the second equality holds because g is symmetric. Hence, for $\beta \in (0, \beta_A]$,

$$q^w g'(\lambda_2 h + \beta) + (1 - q^w)g'(-\lambda_2 h + \beta) < 0$$

and therefore

$$q^w g(\lambda_2 h + \beta_A) + (1 - q^w)g(-\lambda_2 h + \beta_A) \leq q^w g(\lambda_2 h + \beta_B) + (1 - q^w)g(-\lambda_2 h + \beta_B), \quad (52)$$

with strict inequality if $\beta_B < \beta_A$. Since $\beta_B > 0$ and, under the hypothesis, $q^l \leq q^w$, the right-hand side of (52) is less than or equal to $q^l g(\lambda_2 h + \beta_B) + (1 - q^l)g(-\lambda_2 h + \beta_B)$. Hence (50) and (51) imply that $C'_2(e_2^w) \leq C'_2(e_2^l)$, and by the convexity of C_2 it follows that $e_2^w \leq e_2^l$. Since in equilibrium, Δe_1 must have the same sign as $e_2^l - e_2^w \geq 0$, we obtain a contradiction to our assumption that $\Delta e_1 \leq -\alpha < 0$.

Similarly, if $\Delta e_1 \geq 0$, then it follows from $\alpha + \Delta e_1 > 0$ that $q^l > q^w$, so $\beta_B > \beta_A$. Now we have, for all $\beta \in (0, \beta_B)$,

$$\frac{q^l}{1 - q^l} = \frac{g(\lambda_2 h - \beta_B)}{g(\lambda_2 h + \beta_B)} > \frac{g(\lambda_2 h - \beta)}{g(\lambda_2 h + \beta)} > \frac{g'(\lambda_2 h - \beta)}{g'(\lambda_2 h + \beta)} = -\frac{g'(-\lambda_2 h + \beta)}{g'(\lambda_2 h + \beta)},$$

and thus

$$q^l g'(\lambda_2 h + \beta) + (1 - q^l)g'(-\lambda_2 h + \beta) < 0.$$

Hence, since $\beta_B > \beta_A > 0$,

$$q^l g(\lambda_2 h + \beta_B) + (1 - q^l) g(-\lambda_2 h + \beta_B) < q^l g(\lambda_2 h + \beta_A) + (1 - q^l) g(-\lambda_2 h + \beta_A),$$

and the right-hand side is strictly smaller than $q^w g(\lambda_2 h + \beta_A) + (1 - q^w) g(-\lambda_2 h + \beta_A)$ because $q^l > q^w$. So it follows from (50) and (51) that $C'_2(e_2^w) > C'_2(e_2^l)$ and thus $e_2^w > e_2^l$. Since in equilibrium, Δe_1 must have the same sign as $e_2^l - e_2^w < 0$, we obtain a contradiction to our assumption that $\Delta e_1 \geq 0$.

Part (ii) Let $(\beta^{II}, \Delta e_1^{II})$ and $(\beta_A^{ID}, \beta_B^{ID}, \Delta e_1^{ID})$ denote the unique equilibrium with identity-independent and identity-dependent biases, respectively. Assume that h is sufficiently small such that $-\lambda_2 h + \beta_A^{ID} \geq 0$. Choosing h like that is possible because, by analogy to Proposition 1, it holds that $\lim_{h \rightarrow 0} \beta_A^{ID} > 0$. We now show that

$$\Delta e_1^{ID} < \Delta e_1^{II}. \quad (53)$$

By contradiction, assume that $\Delta e_1^{ID} \geq \Delta e_1^{II}$. Starting from $(\beta_A^{ID}, \beta_B^{ID}, \Delta e_1^{ID})$ suppose the principal is restricted to use identity-independent bias, resulting in the choice $\hat{\beta} = \beta^*(\alpha + \Delta e_1^{ID})$. Consider the agents' corresponding effort response $\Delta e_1^*(\hat{\beta}, \hat{\beta})$. As costs are quadratic it follows from (47) that the agents' first-stage effort differential satisfies the implicit equation

$$c_1 \Delta e_1 - [C_2(e_2^l) - C_2(e_2^w)] \sum_{\Delta a \in \{-h, h\}} g(\lambda_1 \Delta a + \alpha + \Delta e_1) = 0, \quad (54)$$

where, using (50) and (51),

$$\begin{aligned} C_2(e_2^l) - C_2(e_2^w) &= \frac{1}{c_2} [q^l g(-\lambda_2 h - \beta_B) + (1 - q^l) g(\lambda_2 h - \beta_B)]^2 \\ &\quad - \frac{1}{c_2} [q^w g(\lambda_2 h + \beta_A) + (1 - q^w) g(-\lambda_2 h + \beta_A)]^2. \end{aligned}$$

Because $\beta_A^{ID} < \beta^*(\alpha + \Delta e_1^{ID}) < \beta_B^{ID}$ as shown above, the move from $\beta_A = \beta_A^{ID}$ and $\beta_B = \beta_B^{ID}$ to $\beta_A = \beta_B = \beta^*(\alpha + \Delta e_1^{ID})$ decreases $g(\lambda_2 h + \beta_A)$ and increases $g(-\lambda_2 h - \beta_B)$ and, given $-\lambda_2 h + \beta_A^{ID} \geq 0$ (which implies $\lambda_2 h - \beta_B^{ID} < 0$) it also decreases $g(-\lambda_2 h + \beta_A)$ and increases $g(\lambda_2 h - \beta_B)$. The move from $(\beta_A^{ID}, \beta_B^{ID})$ to $\beta^*(\alpha + \Delta e_1^{ID})$ thus reduces (54) for any fixed Δe_1 by increasing $C_2(e_2^l) - C_2(e_2^w)$, which is negative, as shown in the proof of claim (i). Given that (54) is negative for $\Delta e_1 = -\alpha$ and positive for $\Delta e_1 = 0$ and equilibrium is unique (which is guaranteed by the assumption that c_1 is sufficiently large), the move from $(\beta_A^{ID}, \beta_B^{ID})$ to $\beta^*(\alpha + \Delta e_1^{ID})$ thus leads to an increase in Δe_1 , i.e. we have

shown that $\Delta e_1^*(\beta^*(\alpha + \Delta e_1^{ID}), \beta^*(\alpha + \Delta e_1^{ID})) > \Delta e_1^*(\beta_A^{ID}, \beta_B^{ID}) = \Delta e_1^{ID}$.

To see that this leads to a contradiction, let $\gamma = (\beta^*)^{-1} - \alpha$. Then $\gamma(\beta)$ gives the conjectured effort differential Δe_1 that makes β the principal's optimal choice of identity-independent bias. Given uniqueness of the equilibrium $(\beta^{II}, \Delta e_1^{II})$, the curves $\gamma(\beta)$ and $\Delta e_1^*(\beta, \beta)$ intersect exactly once. And because $\Delta e_1^*(\beta, \beta)$ goes to zero for $\beta \rightarrow 0$ and for $\beta \rightarrow \infty$ and $\gamma(\beta)$ is strictly decreasing, $\Delta e_1^*(\beta, \beta)$ must cross $\gamma(\beta)$ from below. In particular, for any $\beta < \beta^{II}$ it must hold that $\gamma(\beta) > \Delta e_1^*(\beta, \beta)$. Note that $\hat{\beta} = \beta^*(\alpha + \Delta e_1^{ID}) < \beta^*(\alpha + \Delta e_1^{II}) = \beta^{II}$ because β^* is decreasing and we have assumed that $\Delta e_1^{ID} > \Delta e_1^{II}$. Hence $\gamma(\hat{\beta}) > \Delta e_1^*(\hat{\beta}, \hat{\beta})$, or formulated equivalently, $\Delta e_1^{ID} = \Delta e_1^*(\beta_A^{ID}, \beta_B^{ID}) > \Delta e_1^*(\beta^*(\alpha + \Delta e_1^{ID}), \beta^*(\alpha + \Delta e_1^{ID}))$, which contradicts our earlier finding. ■

Proof of Corollary 4

This proof assumes that α is sufficiently small such that $P_\alpha(\beta^*(\tilde{\alpha}), \beta^*(\tilde{\alpha}), \tilde{\alpha})$ is increasing in $\tilde{\alpha}$ for all $\tilde{\alpha} < \alpha + \Delta e_1^{II}$. Such values of α exist because $P_\alpha(\beta^*(\tilde{\alpha}), \beta^*(\tilde{\alpha}), \tilde{\alpha})$ is increasing in $\tilde{\alpha}$ for small $\tilde{\alpha}$ since $\frac{\partial P_\alpha}{\partial \tilde{\alpha}} > 0$ and, as shown in the proof of Proposition 4, $\lim_{\alpha \rightarrow 0} \Delta e_1^{II} = 0$ and $\lim_{\tilde{\alpha} \rightarrow 0} \frac{\partial \beta^*}{\partial \tilde{\alpha}} = 0$.

Part (i) This claim is true because $P_\alpha(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) < P_\alpha(\beta^{II}, \beta^{II}, \alpha + \Delta e_1^{II})$ follows from (53). To see this, note first that, as P_α is increasing in β_A but decreasing in β_B , it holds that $P_\alpha(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) < P_\alpha(\beta^*(\alpha + \Delta e_1^{ID}), \beta^*(\alpha + \Delta e_1^{ID}), \alpha + \Delta e_1^{ID})$, because, as shown above, the principal's optimal biases satisfy $\beta_A^{ID} = \beta_A^*(\alpha + \Delta e_1^{ID}) < \beta^*(\alpha + \Delta e_1^{ID}) < \beta_B^*(\alpha + \Delta e_1^{ID}) = \beta_B^{ID}$. And because, by assumption, $P_\alpha(\beta^*(\tilde{\alpha}), \beta^*(\tilde{\alpha}), \tilde{\alpha})$ is increasing in $\tilde{\alpha}$ for all $\tilde{\alpha} < \alpha + \Delta e_1^{II}$ it follows from (53) that $P_\alpha(\beta^*(\alpha + \Delta e_1^{ID}), \beta^*(\alpha + \Delta e_1^{ID}), \alpha + \Delta e_1^{ID}) < P_\alpha(\beta^*(\alpha + \Delta e_1^{II}), \beta^*(\alpha + \Delta e_1^{II}), \alpha + \Delta e_1^{II}) = P_\alpha(\beta^{II}, \beta^{II}, \alpha + \Delta e_1^{II})$.

Part (ii) To prove the second claim, we determine the effect of a move from identity-independent to identity-dependent bias on the agents' utility differential by considering

$$\lim_{\alpha \rightarrow 0} \frac{d}{d\alpha} \Delta U(\beta_A, \beta_B, \alpha + \Delta e_1) = 2 \frac{dP_\alpha}{d\alpha} \Big|_{\alpha=0} + c_1 e_1^* \frac{\partial \Delta e_1}{\partial \alpha} \Big|_{\alpha=0}.$$

Here we used that costs are quadratic and that in the limit as $\alpha \rightarrow 0$ agents exert the same first-stage effort $e_1^* = \lim_{\alpha \rightarrow 0} e_{A,1}^* = \lim_{\alpha \rightarrow 0} e_{B,1}^*$. When $\alpha \rightarrow 0$, the first-order conditions (45) and (46) both simplify to

$$c_1 e_1^* = g(\lambda_1 h) [G(\lambda_2 h + \beta^*(0)) - G(\lambda_2 h - \beta^*(0))].$$

Using the fact that, as shown in the proof of Proposition 4, $\lim_{\alpha \rightarrow 0} \frac{\partial \beta_A^*}{\partial \alpha} = -\lim_{\alpha \rightarrow 0} \frac{\partial \beta_B^*}{\partial \alpha}$,

we get

$$\begin{aligned} \frac{dP_\alpha^{ID}}{d\alpha}\Big|_{\alpha=0} &= g(\lambda_1 h)[G(\lambda_2 h + \beta^*(0)) - G(\lambda_2 h - \beta^*(0))] \left(1 + \frac{d\Delta e_1^{ID}}{d\alpha}\Big|_{\alpha=0}\right) \\ &\quad + [G(\lambda_1 h)g(\lambda_2 h + \beta^*(0)) + G(-\lambda_1 h)g(-\lambda_2 h + \beta^*(0))] \frac{d\beta_A^*}{d\alpha}\Big|_{\alpha=0}, \end{aligned}$$

whereas $\lim_{\alpha \rightarrow 0} \frac{\partial \beta^*}{\partial \alpha} = 0$ implies

$$\frac{dP_\alpha^{II}}{d\alpha}\Big|_{\alpha=0} = g(\lambda_1 h)[G(\lambda_2 h + \beta^*(0)) - G(\lambda_2 h - \beta^*(0))] \left(1 + \frac{d\Delta e_1^{II}}{d\alpha}\Big|_{\alpha=0}\right).$$

For the difference we thus get

$$\begin{aligned} \lim_{\alpha \rightarrow 0} \frac{d(\Delta U^{ID} - \Delta U^{II})}{d\alpha} &= g(\lambda_1 h)[G(\lambda_2 h + \beta^*(0)) - G(\lambda_2 h - \beta^*(0))] \frac{d(\Delta e_1^{ID} - \Delta e_1^{II})}{d\alpha}\Big|_{\alpha=0} \\ &\quad + [G(\lambda_1 h)g(\lambda_2 h + \beta^*(0)) + G(-\lambda_1 h)g(-\lambda_2 h + \beta^*(0))] \frac{d\beta_A^*}{d\alpha}\Big|_{\alpha=0}. \end{aligned}$$

This is strictly negative, because $\frac{d\beta_A^*}{d\alpha}\Big|_{\alpha=0} < 0$ as shown in the proof of Proposition 4 and because our analysis above implies that $\Delta e_1^{ID} - \Delta e_1^{II}$ must be non-increasing for small α . Given that for $\alpha \rightarrow 0$, $\Delta U^{ID} = \Delta U^{II} = 0$, for small α it must therefore hold that $\Delta U(\beta_A^{ID}, \beta_B^{ID}, \alpha + \Delta e_1^{ID}) < \Delta U(\beta_A^{II}, \beta_B^{II}, \alpha + \Delta e_1^{II})$. ■

References

- Alesina, Alberto and George Marios Angeletos**, “Fairness and redistribution,” *American Economic Review*, 2005, 95 (4), 960–980.
- , **Stefanie Stantcheva, and Edoardo Teso**, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, 2018, 108 (2), 521–554.
- Almås, Ingvild, Alexander W. Cappelen, and Bertil Tungodden**, “Cutthroat capitalism versus cuddly socialism: Are americans more meritocratic and efficiency-seeking than Scandinavians?,” *Journal of Political Economy*, 2020, 128 (5), 1753–1788.
- Altonji, Joseph G. and Charles R. Pierret**, “Employer learning and statistical discrimination,” *Quarterly Journal of Economics*, 2001, 116 (1), 313–350.
- Ariga, Kenn, Yasushi Ohkusa, and Giorgio Brunello**, “Fast track: Is it in the genes? The promotion policy of a large Japanese firm,” *Journal of Economic Behavior and Organization*, 1999, 38 (4), 385–402.

- Azmat, Ghazala and Anne Boring**, “Gender diversity in firms,” *Oxford Review of Economic Policy*, Winter 2020, 36 (4), 760–782.
- Baker, George, Michael Gibbs, and Bengt Holmström**, “The Internal Economics of the Firm: Evidence from Personnel Data,” *Quarterly Journal of Economics*, 1994, 109 (4), 881–919.
- Bardhi, Arjada, Yingni Guo, and Bruno Strulovici**, “Early-Career Discrimination: Spiraling or Self-Correcting?,” 2023. https://yingniguo.com/wp-content/uploads/2023/04/draft_39.pdf.
- Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka**, “Attention discrimination: Theory and field experiments with monitoring information acquisition,” *American Economic Review*, 2016, 106 (6), 1437–1475.
- Bedard, Kelly and Elizabeth Dhuey**, “The Persistence of Early Childhood Maturity: International Evidence of Long-Run Age Effects,” *The Quarterly Journal of Economics*, 2006, 121 (4), 1437–1472.
- Belzil, Christian and Michael Bognanno**, “Promotions, demotions, halo effects, and the earnings dynamics of American executives,” *Journal of Labor Economics*, 2008, 26 (2), 287–310.
- Benabou, Roland and Jean Tirole**, “Belief in a Just World and Redistributive Politics,” *Quarterly Journal of Economics*, 2006, 121 (2), 699–746.
- Benson, Alan M., Danielle Li, and Kelly Shue**, “Potential and the Gender Promotion Gap,” *Academy of Management Proceedings*, 2023, 2023 (1).
- Bergemann, Dirk and Juuso Välimäki**, “Bandit Problems,” *The New Palgrave Dictionary of Economics*, 2008, pp. 1–7.
- Blank, Rebecca M.**, “Tracing the Economic Impact of Cumulative Discrimination,” *American Economic Review*, 2005, 95 (2), 99–103.
- Bohren, Aislinn J., Peter Hull, and Alex Imas**, “Systemic Discrimination: Theory and Measurement,” 2023. https://www.aislinnbohren.com/s/Systemic_Discrimination_Dec22_2023.pdf.
- Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg**, “The dynamics of discrimination: Theory and evidence,” *American Economic Review*, 2019, 109 (10), 3395–3436.

- Bukodi, Erzsébet, John H. Goldthorpe, and Inga Steinberg**, “Social stratification in science: the ultra-elite in the UK,” *European Societies*, 2024, pp. 1–42.
- Cappelen, Alexander W., Astri Drange Hole, Erik Ø. Sørensen, and Bertil Tungodden**, “The Pluralism of Fairness Ideals: An Experimental Approach,” *American Economic Review*, 2007, *97* (3), 818–827.
- , **James Konow, Erik Ø. Sørensen, and Bertil Tungodden**, “Just luck: An experimental study of risk-taking and fairness,” *American Economic Review*, 2013, *103* (4), 1398–1413.
- Ciocca Eller, Christina**, “What Makes a Quality College? Reexamining the Equalizing Potential of Higher Education in the United States,” *American Journal of Sociology*, 2023, *129* (3), 637–714.
- Coate, Stephen and Glenn Loury**, “Will Affirmative-Action Policies Eliminate Negative Stereotypes?,” *American Economic Review*, 1993, *83* (5), 1220–1240.
- Cole, Stephen, Jonathan R. Cole, and Gary A. Simon**, “Chance and Consensus in Peer Review,” *Science*, 1981, *214* (4523), 881–886.
- Cong, Lin William and Yizhou Xiao**, “Persistent Blessings of Luck: Theory and an Application to Venture Capital,” *Review of Financial Studies*, 2022, *35* (3), 1183–1221.
- Deaner, Robert O., Aaron Lowen, and Stephen Cobley**, “Born at the Wrong Time: Selection Bias in the NHL Draft,” *PLoS ONE*, 2013, *8* (2).
- DiPrete, Thomas A. and Gregory M. Eirich**, “Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments,” *Annual Review of Sociology*, 2006, *32*, 271–297.
- Du, Qianqian, Huasheng Gao, and Maurice D. Levi**, “The relative-age effect and career success: Evidence from corporate CEOs,” *Economics Letters*, 2012, *117* (3), 660–662.
- Duclos, Jean-Yves, Joan Esteban, and Debraj Ray**, “Polarization: Concepts, Measurement, Estimation,” *Econometrica*, 2004, *72* (6), 1737–1772.
- Durandard, Théo**, “Dynamic delegation in promotion contests,” 2023. <https://arxiv.org/abs/2308.05668>.

- Exley, Christine and Kirby Nielsen**, “The Gender Gap in Confidence: Expected But Not Accounted For,” *American Economic Review*, 2024, 114 (3), 851–885.
- Farber, Henry and Robert Gibbons**, “Learning and Wage Dynamics,” *Quarterly Journal of Economics*, 1996, 111 (4), 1007–1047.
- Fong, Christina**, “Social preferences, self-interest, and the demand for redistribution,” *Journal of Public Economics*, 2001, 82, 225–246.
- Forbes, J. Benjamin**, “Early Intraorganizational Mobility: Patterns and Influences,” *Academy of Management Journal*, 1987, 30 (1), 110–125.
- Fosgerau, Mogens, Rajiv Sethi, and Jörgen Weibull**, “Equilibrium Screening and Categorical Inequality,” *American Economic Journal: Microeconomics*, 2023, 15 (3), 201–242.
- Frank, Robert H.**, *Success and Luck: Good Fortune and the Myth of Meritocracy*, Princeton University Press, 2016.
- Fryer, Roland G.**, “Belief flipping in a dynamic model of statistical discrimination,” *Journal of Public Economics*, 2007, 91 (5-6), 1151–1166.
- Gamoran, Adam and Robert Mare**, “Secondary school tracking and educational inequality: compensation, reinforcement or neutrality? American Journal of Sociology,” *American Journal of Sociology*, 1989, 94 (5), 1146–1183.
- Gibbons, Robert and Michael Waldman**, “A Theory of Wage and Promotion Dynamics inside Firms,” *Quarterly Journal of Economics*, 1999, 114 (4), 1321–1358.
- **and** –, “Enriching a theory of wage and promotion dynamics inside firms,” *Journal of Labor Economics*, 2006, 24 (1), 59–107.
- Gill, David and Victoria Prowse**, “Gender differences and dynamics in competition: The role of luck,” *Quantitative Economics*, 2014, 5 (2), 351–376.
- Kahn, Lisa B. and Fabian Lange**, “Employer learning, productivity, and the earnings distribution: Evidence from performance measures,” *Review of Economic Studies*, 2014, 81 (4), 1575–1613.
- Konow, James**, “Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions,” *American Economic Review*, 2000, 90 (4), 1072–1092.

- Lange, Fabian**, “The speed of employer learning,” *Journal of Labor Economics*, 2007, 25 (1), 1–35.
- Lavy, Victor and Rigissa Megalokonomou**, “The Short- and the Long-Run Impact of Gender-Biased Teachers,” *American Economic Journal: Applied Economics*, 2024, 16 (2), 176–218.
- Lazear, Edward P.**, “Performance pay and productivity,” *American Economic Review*, 2000, 90 (5), 1346–1361.
- , “Compensation and Incentives in the Workplace,” *Journal of Economic Perspectives*, 2018, 32 (3), 195–214.
- **and Sherwin Rosen**, “Rank-order tournaments as optimum labor contracts,” *Journal of Political Economy*, 1981, 89 (5), 841–864.
- Lefgren, Lars J., David P. Sims, and Olga B. Stoddard**, “Effort, luck, and voting for redistribution,” *Journal of Public Economics*, 2016, 143, 89–97.
- Lepage, Louis-Pierre**, “Experience-based Discrimination,” *American Economic Journal: Applied Economics*, 2024.
- Li, Danielle, Lindsey Raymond, and Peter Bergman**, “Hiring as Exploration,” 2024. <https://danielle-li.github.io/assets/docs/HiringAsExploration.pdf>.
- Lluis, Stéphanie**, “The Role of Comparative Advantage and Learning in Wage Dynamics and Intrafirm Mobility: Evidence from Germany,” *Journal of Labor Economics*, 2005, 23 (4), 725–767.
- Lundberg, Shelly J. and Richard Startz**, “Private Discrimination and Social Intervention in Competitive Labor Market,” *American Economic Review*, 1983, 73 (3), 340–347.
- Merton, Robert K.**, “The Matthew effect in science,” *Science*, 1968, 159 (3810), 56–62.
- Meyer, Margaret A.**, “Learning from coarse information: Biased contests and career profiles,” *Review of Economic Studies*, 1991, 58 (1), 15–41.
- Moro, Andrea and Peter Norman**, “A general equilibrium model of statistical discrimination,” *Journal of Economic Theory*, 2004, 114 (1), 1–30.

- Nanda, Ramana, Sampsa Samila, and Olav Sorenson**, “The persistent effect of initial success: Evidence from venture capital,” *Journal of Financial Economics*, 2020, *137* (1), 231–248.
- Onuchic, Paula**, “Recent Contributions to Theories of Discrimination,” 2023. <http://arxiv.org/abs/2205.05994>.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz**, “The short- and long-term career effects of graduating in a recession,” *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.
- Pastorino, Elena**, “Careers in Firms: The Role of Learning about Ability and Human Capital Acquisition,” *Journal of Political Economy*, 2024, *132* (6), 1994–2073.
- Piketty, Thomas**, *Capital in the Twenty-First Century*, Cambridge, MA: The Belknap Press of Harvard University Press, 2014.
- Roemer, John E.**, “Equality of Opportunity,” in Kenneth Arrow, Samuel Bowles, and Steven Durlauf, eds., *Meritocracy and Economic Inequality*, Princeton (NJ): Princeton University Press, 2000, pp. 17–32.
- Rosenbaum, James E.**, “Tournament Mobility: Career Patterns in a Corporation,” *Administrative Science Quarterly*, 1979, *24* (2), 220–241.
- Sandel, Michael**, *The Tyranny of Merit*, New York: Farrar, Straus and Giroux, 2020.
- Sen, Amartya**, “Merit and Justice,” in Kenneth Arrow, Samuel Bowles, and Steven Durlauf, eds., *Meritocracy and Economic Inequality*, Princeton (NJ): Princeton University Press, 2000, pp. 5–16.
- Sethi, Rajiv and Rohini Somanathan**, “Meritocracy and Representation,” *Journal of Economic Literature*, 2023, *61*, 941–957.
- Young, Michael**, *The Rise of Meritocracy 1870-2033*, London: Thames and Hudson, 1958.
- Zhang, Sam, K. Hunter Wapman, Daniel B. Larremore, and Aaron Clauset**, “Labor advantages drive the greater productivity of faculty at elite universities,” *Science Advances*, 2022, *8* (46), 1–22.