Selecting the Best: The Persistent Effects of Luck*

Mikhail Drugov,† Margaret Meyer‡ and Marc Möller§

January 23, 2024

Abstract

Meritocratic principles seem abandoned when an initial stroke of luck affects the final allocation of economic resources or decision-making authority. This article proposes a stylized model of organizational learning where agents' performance at each stage is the sum of time-invariant, unobservable ability, privately-chosen effort, and transitory noise. Our main result shows that to identify the most able agent, selection has to be biased in favor of agents who perform well initially, even in the limit when noise swamps ability- and effort-differentials in the determination of outcomes. Making luck persistent is thus rationalized as a necessary feature of selective efficiency. The role of luck becomes amplified by strategic behavior of informed agents and towards the top of an organization's hierarchy, where performance measurement is limited to be ordinal rather than cardinal.

Keywords: Organizational learning; Selective efficiency; Bias; Promotions.
JEL *classification*: D21, D82, D83, M51.

*We thank Paul Klemperer, Dmitry Ryvkin and participants of the 2021 Transatlantic Theory Workshop, the 8th Annual Conference on "Contests: Theory and Evidence", the 2022 Conference on Mechanism and Institution Design, and the 2023 European Summer Symposium in Economic Theory for valuable comments. All errors are our own.

†New Economic School 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@vwi.unibe.ch.

1 Introduction

Sometimes an individual's success is explained, or even discredited, as resulting from an initial stroke of good luck. Bestselling authors such as Gladwell (2008) or Frank (2016) document a multitude of careers of over-achievers, ranging from the arts to business, that were kick-started by fortunate circumstances or events. Even major scientific achievements, such as special relativity theory, are debated to be the consequence of external factors rather than genius (Gallison, 2003). The idea that, in the extreme, success is no longer attributable to skill-differentials finds support in empirical evidence (Keuschnigg et al., 2023) and numerical simulations (Denrell and Liu, 2012), emphasizing the importance of the serial correlation of outcomes.

Underlying such narratives is the concern that economic institutions or organizational practices may amplify the role of luck by making its effects long-lasting. A common argument, across different disciplines of the Social Sciences, is that resources, favoritism, or *biases* granted to early performers will increase the likelihood with which 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 (Rosenbaum, 1979; Forbes, 1987; Baker et al., 1994) inflate the importance of early performance for final success. As a consequence of such policies, chance events such as graduating during a recession or being the oldest kid in class can have long-lasting effects on both, labor market outcomes (Oreopoulos et al., 2012) and educational achievements (Bedard and Dhuey, 2006).[1] A related phenomenon in Finance, known as *rich-get-richer dynamics*, describes the finding that "being at the right time in the right place" creates future investment opportunities capable of explaining performance persistence of hedge funds (Cong and Xiao, 2022) and venture capitalists (Nanda et al., 2020). Finally, the Sociology literature uses the terms *Matthew principle* or *cumulative advantage* (Merton, 1968; DiPrete and Eirich, 2006) and argues that performance differentials, such as outstanding 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).

By implementing biases, the aforementioned institutions induce correlation between initial and final outcomes in situations where noise has a substantial influence on performance. They hereby improve *selective efficiency*, i.e. the allocation of resources to the most productive individuals, as long as initial success can be attributed to *merit*, commonly defined as a combination of individual talents or actions (Sen, 2000). However, in

[1]Evidence for *relative age effects* has been documented for physical activities, such as becoming a player in the National Hockey League (Deaner et al., 2013), as well as intellectual achievements such as being selected as a CEO of a S&P 500 company (Du et al., 2012).

the limit where noise swamps ability and effort in the determination of outcomes, biases merely *make luck persistent* by inducing final outcomes to depend on initial conditions that are entirely random. In this paper we argue that, while seemingly at odds with meritocratic principles, making luck persistent can be rationalized as a characteristic feature of an organization's "selection of the best".[2]

By explaining how institutions shape the role of luck for individual success, our theory contributes to our understanding of the "origin" of economic inequality. This is important because economic 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). 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). The connection between luck and success also affects 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), individual experiences concerning the role of luck for success may contribute to shape political outcomes. This is especially relevant when beliefs about luck select 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).[3] By rationalizing the persistent effects of luck as an institutional practice and by investigating its determinants, our theory may thus contribute to our understanding of cross-country differences in economic inequality and political polarization.

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 relative abilities. The organization observes only the ordinal ranking of performances at each stage and attempts to select the most able agent.[4] Agents choose efforts to maximize the probability

[2]The term "meritocracy" originates from Young's (1958) apocalyptic vision a future society in which "merit" serves as the central determinant of an individual's power and wealth. In spite of a dispute over what constitutes merit, modern democracies claim to have adopted merit as a basis for their allocation of resources and decision-making power (Piketty, 2014).

[3]Experimental 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).

[4]Ordinal performance measurement arises naturally when performance is difficult to quantify. For instance, 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. Arguably, ordinal performance measurement is thus the most relevant case to consider in situations where selection matters most, i.e. towards the top of an organization's hierarchy.

of being selected, minus their effort costs.[5] The organization's optimal selection rule augments the second-stage performance of the agent who performed best in the first stage by use of an additive bias and selects the agent who performs best in the second stage. Our main variable of interest is the persistence of early success, i.e. the probability with which, in equilibrium, the agent with the better initial performance becomes selected in the final stage.

We start our analysis in Section 3 by showing that, when agents have no private information about their relative abilities, effort choices will be identical across agents in *both* stages, in spite of the asymmetries induced by learning from first-stage performance and application of second-stage bias. In the absence of private information, effort choices thus have no effect on selective efficiency, implemented bias, nor persistence. Our main result shows that in the limit, where 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, optimal bias makes first-stage "winners" considerably more likely to become selected than first-stage "losers". This shows that the persistence of luck illustrated by our motivating examples does not have to be a consequence of too much or the wrong kind of bias being employed, but can be explained as an institutionalized feature of an organization's attempt to "select the best" in environments where individual performance is noisy.

The basic intuition for this result can be explained as follows. Optimal bias balances the informativeness of the agents' performance rankings across stages and is such that, if the first-stage loser just managed to achieve a hypothetical – because unobservable – second-stage draw despite the bias against him, the organization would be indifferent as to which agent to select. In the limit, both, an unbiased first-stage win and a second-stage draw against bias become uninformative which means that optimal bias has to equate the *rates* at which their informativeness varies with the agents' heterogeneity. However, setting bias equal to zero not only nullifies the informativeness of a second-stage draw but also the rate at which this informativeness changes with the agents' heterogeneity. In other words, a strictly positive bias arises 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.

In Section 4 we consider how the persistence of luck resulting from an organization's use of bias is affected by informed agents' strategic behavior by turning attention to

[5] Lazear and Rosen (1981) argue that competing to become selected, e.g. for promotion, can provide workers inside firms with efficient incentives to exert effort and may thus substitute incentive schemes that rely on cardinal performance measurement when performance is hard to quantify.

the case where agents share private information about their relative abilities. Because, in our setting, effort acts as a substitute for ability in creating 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 a second-stage bias. We show that, contrary to this intuition, the agent thought more likely to be the more able agent exerts a strictly larger first-stage effort than his rival. Effort choices thus reinforce the agents' ability differential and under mild conditions on the agents' cost functions, equilibrium bias can be shown to be increasing in the agents' informational advantage. Hence, while biased selection arises from the organization's ignorance of the agents' abilities, information on the side of the agents does not mitigate but rather amplifies the persistent effects of luck. This result resonates well with the dominant 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 obtain an informational advantage over their superiors regarding their co-workers abilities.

To provide further insight into the relation between selective efficiency and the persistence of luck, Section 5 considers the counter-factual situation where performance information is *cardinal* rather than ordinal. In this case, the optimal bias can condition on the first-stage margin of victory. We show that in the limit, the equilibrium bias under ordinal information equals an appropriately-defined average of the optimal cardinal biases, as the margin of victory varies. We provide a sufficient condition under which a restriction to ordinal rather than cardinal performance information makes luck more persistent (on average). This condition is satisfied when the distribution of noise is "sufficiently normal" and our result then implies that organizations can be expected to make luck more persistent when individual performance can be ranked but is hard to quantify. As ordinal performance measurement is more prevalent towards the top of an organization's hierarchy, our theory thus highlights the special role of luck for selection into those positions where the choice of the "right" or the "wrong" agent can be most consequential.

Related literature. Our paper contributes to the literature on organizational learning. Driven by emerging 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 within firms. A robust empirical finding is that early wage increases and early promotions increase the probability to become promoted at a later stage. Whether this correlation is caused by worker heterogeneity 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). While in the seminal models serial correlation of promotion rates arises from workers' time-invariant ability differences, our model shows that even when ability differences become negligible, 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".[6] 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 a task assignment augments the persistence of luck because a greater bias is required to raise the informativeness of the less informative later tasks. Her structural estimates provide strong evidence that learning, besides human capital accumulation, has a sizeable impact on career outcomes. Finally, the "pattern of promotions" that arises from optimal biased selection is reminiscent of the "late beginner's effect" documented by Chiappori et al. (1999) which asserts that workers who perform well after performing badly have a higher chance to become promoted than workers with performances in reversed order. While their explanation relies on the influence of wage rigidity on symmetric employer learning, in our model "late beginners" become selected because their late win against a bias is a more informative signal about their ability than their early loss in the absence of bias.

2 Model

We consider an organization consisting of a risk-neutral principal and two agents $i \in \{A, B\}$ with potentially heterogeneous abilities. The principal needs to select one of the agents for a future task (e.g. CEO or chief of staff) whose payoff to the principal, $\Pi(a)$, is increasing in the selected agent's ability, a. The principal can base his selection on agents' performance but his choice is complicated by the fact that neither the agents' abilities nor their efforts are observable and that performance information is noisy and only ordinal rather than cardinal.

Agents' abilities $a_i \in \Re$ are assumed to be distributed on $\{\underline{a}, \underline{a} + h\}$ with parameter

[6]Using 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, Lluis (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).

$h > 0$ denoting their potential *heterogeneity*.[7] The principal and the agents share a common prior, $q^0 \equiv \mathbb{P}(a_A > a_B | a_A \neq a_B)$, but for the principal agents A and B are indistinguishable. This allows us to differentiate between the case $q^0 \in (\frac{1}{2}, 1]$ where agents have superior information about their relative abilities and the case $q^0 = \frac{1}{2}$ where agents are equally uninformed as the principal.[8]

To capture the dynamic nature of organizational learning, we assume that the principal observes the agents' relative performance during two stages. In each stage $t \in \{1, 2\}$, agent i's performance, $x_{i,t} \in \Re$, 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 $\epsilon_{i,t} \in \Re$.[9] That is,

$$x_{i,t} \equiv \lambda_t a_i + e_{i,t} + \epsilon_{i,t}. \tag{1}$$

Variation in λ_t across stages accounts for potential differences in the impact of ability on performance. This is especially relevant when agents' task changes over time, for instance, because firms' choose first assignments for their managers that are particularly informative about their abilities, as argued by Pastorino (2024). We assume that the principal can identify the agent with the higher performance $x_{i,1}$ as the winner of the first stage, with ties broken randomly. In the second stage, the principal may assign a bias $\beta \in \Re$ to the winner of the first stage. Having won the first stage, agent i is then identified as the winner of the second stage if $x_{i,2} + \beta > x_{j,2}$.

The principal's objective is to choose the size of the bias β to maximize his payoff Π. Given our assumptions, this objective is equivalent to maximizing *selective efficiency*, $S(\beta; h)$, defined as the probability that, conditional on agents differing in their abilities, the more able agent becomes selected. When bias is chosen optimally, selecting the winner of the second stage maximizes the principal's objective.[10] Note that we assume that the principal chooses the size of the bias β *after* observing the agents' first-stage ranking rather than committing to it upfront. This assumption is motivated by the observation that in many instances bias is implemented in the form of favoritism, making the absence

[7] Our results do not rely on ability having binary support and remain valid when $a_i = \underline{a} + h\alpha_i$, where the joint distribution of $(\alpha_i, \alpha_j) \in \Re^2$ is symmetric with respect to the two components but otherwise arbitrary.

[8] Virtually all the employer learning models reviewed in the Introduction assume that workers are ignorant about their own ability, and hence correspond to the case where $q^0 = \frac{1}{2}$.

[9] By logarithmic transformation, our results remain qualitatively unchanged when performance equals the product rather than the sum of ability, effort, and noise.

[10] A randomly assigned first-stage bias does not increase selective efficiency, as long as the size of the second-stage bias cannot condition on the first-stage winner's *identity*, $i \in \{A, B\}$, which is consistent with our assumption that agents are indistinguishable for the principal. We thus abstract from the possibility that bias is assigned in *both* stages.

of such commitment the, arguably, more relevant case.[11]

Agents choose their efforts in each stage to maximize the difference between their chances to become selected and their effort costs. Effort costs are assumed to be identical across agents but may differ across stages. More specifically, we let $C_t(e_{i,t})$ denote agent i's effort cost in stage t and assume that C_t is strictly convex. In equilibrium, agents form correct expectations about the principal's choice of bias.

Because outcomes depend only on performance *differentials*, the distribution of the difference in the individual noise terms, $\Delta\epsilon_t \equiv \epsilon_{A,t} - \epsilon_{B,t}$ is a key primitive of our model. We assume that the random variables $\Delta\epsilon_t$ 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 around 0; (ii) g is strictly log-concave, i.e.,* $\ln g$ *is strictly concave; (iii) g is twice differentiable on $(-z, z)$; (iv)* $\lim_{y\to z} L(y) = \infty$, *where*

$$L(y) \equiv -\frac{g'(y)}{g(y)}. \tag{2}$$

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

Our main variable of interest is the persistence of outcomes induced by the principal's quest for selective efficiency and the agents' desire to become selected. In particular, our subsequent analysis determines the principal's equilibrium choice of bias, $\beta^*(h)$, and the agents' equilibrium efforts to calculate *persistence*, $P(\beta^*(h); h)$, defined as the probability that the winner of the first stage becomes selected after also winning the second stage. With this objective in mind, it is important to note that the key parameter of our model, $h > 0$, captures the degree of agents' heterogeneity but has a broader interpretation as the *ratio* of agents' heterogeneity to the scale of noise.[12] To shed light on the role of

[11]Most of our results continue to hold when the principal can commit to a bias upfront, that is, before agents choose their efforts in stage one.

[12]To see this, introduce a scaling transformation $\Delta\epsilon_t \to \sigma\Delta\epsilon_t$, with $\sigma > 1$, which makes the difference

initial luck for final outcomes, a large part of our analysis will focus on the limit as $h \to 0$, where the scale of noise becomes large *relative* to the agents' heterogeneity. When, in this limit, $P(\beta^*(h); h)$ turns out to be strictly larger than one half, we will say that *luck is made persistent*, because when $\lim_{h\to 0} P(\beta^*(h); h) > \frac{1}{2}$, the first-stage winner has a greater chance to be selected than the first-stage loser, in spite of the first-stage outcome being entirely random.

3 Persistence of luck

In this section, we start our analysis by considering the case where agents are equally uninformed as the principal about their relative abilities. It turns out that, in this case, the agents' ability to influence their performance through the exertion of effort has no impact on the principal's choice of bias, selective efficiency, nor the persistence of outcomes. This allows us to develop the basic intuition for the connection between these variables, before turning our attention to the effects of informed agents' strategic behavior in the subsequent section.

The following lemma, which we prove in the Appendix, is critical as it shows that, in equilibrium, the efforts of uninformed agents cancel each other in the determination of relative performance.

Lemma 1 (Identical efforts) *Suppose that agents are no better informed about their relative abilities than the principal, i.e.* $q^0 = \frac{1}{2}$. *Then for any anticipated choice of bias 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 proof of Lemma 1, we show that, in the first stage, there exists a pair of identical efforts that are best responses to each other. We then prove that unequal efforts could not be best responses, thereby establishing uniqueness of equilibrium. While in the first stage, equal efforts are a direct consequence of the symmetry imposed by the agents' prior, in the second stage the agent's situation becomes asymmetric, due to the agents' learning from their first-stage ranking and the bias imposed by the principal. We show that in spite of these asymmetries, the agents' marginal returns to effort in the second stage are the same because the pivotal realizations of the difference in noise necessary for winning are identical.

---

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, in equilibrium, the more able agent wins the first stage. It depends on H and σ only through the *heterogeneity-to-noise ratio* $h \equiv \frac{H}{\sigma}$.

Given that, for any level of anticipated bias $\beta \in \Re$, equilibrium efforts cancel at each stage, we can determine the contest's selective efficiency $S(\beta; h)$ easily as follows. Remember that selective efficiency is defined as the probability with which the more able agent wins the contest's final stage. Because the first stage is unbiased, the probability that the more able agent wins the first stage is simply $G(\lambda_1 h) > \frac{1}{2}$. In contrast, for any non-zero value of bias, the probability of winning the second stage depends on the first-stage outcome. If the more able agent won the first stage, then his chance of winning the second stage is $G(\lambda_2 h + \beta)$, whereas if he lost the first stage his chance of winning the second stage is $G(\lambda_2 h - \beta)$. Overall, selective efficiency is thus given by

$$S(\beta; h) = G(\lambda_1 h)G(\lambda_2 h + \beta) + [1 - G(\lambda_1 h)]G(\lambda_2 h - \beta). \tag{3}$$

Differentiating with respect to β and rearranging leads to the following first-order condition for the principal's choice of bias:

$$\frac{G(\lambda_1 h)}{1 - G(\lambda_1 h)} = \frac{g(\lambda_2 h - \beta)}{g(\lambda_2 h + \beta)}. \tag{4}$$

The ratio on the left-hand side is the relative likelihood that a first-stage win identifies the more able agent rather than the less able one. The higher this likelihood ratio, the more informative is the first-stage ranking about agents' relative abilities. The term on the right-hand side is also a likelihood ratio: It is the relative likelihood that a draw in the second stage arises, i.e. $x_{i,2} + \beta = x_{j,2}$, when the more able agent is disadvantaged by the bias compared to when the bias acts in his favor. The higher this likelihood ratio, the more informative would be the hypothetical observation of a second-stage draw about the relative ability of the first-stage loser, who managed to achieve a draw despite being handicapped by the bias. Equation (4) thus shows that optimal bias strikes a balance between the informativeness of the ordinal stage-one ranking – an unbiased win – and the informativeness of the marginal stage-two outcome – a biased draw. In 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 (4) is equal to one and hence strictly smaller than the left-hand side. This is because, without bias, a second-stage draw is uninformative about the agents' abilities. Moreover, given the strict log-concavity of g, as the size of the bias increases, a second-stage draw becomes a strictly stronger signal about the relative ability of the first-stage loser. It thus follows from Assumption 1 that the first-order condition (4) has a unique solution, $\beta^*(h) > 0$, which maximizes selective efficiency. While these arguments suggest that a positive bias will emerge in equilibrium for any level $h > 0$ of agents' heterogeneity, it is not clear what happens in the limit as $h \to 0$. Does bias converge to zero? The following proposition characterizes the value that equilibrium bias takes as the scale of the noise swamps agents heterogeneity in abilities.

Proposition 1 (Equilibrium bias) *Suppose that agents are no better informed about their relative abilities than the principal, i.e. $q^0 = \frac{1}{2}$. The principal's equilibrium choice of bias, $\beta^*(h)$ is strictly positive, even in the limit as noise swamps agents' ability-differences. More specifically, $\beta_0^* \equiv \lim_{h\to 0} \beta^*(h) > 0$ is given by the unique solution of the equation*

$$2\lambda_1 g(0) = \lambda_2 L(\beta_0^*). \tag{5}$$

At first sight, the fact that bias remains strictly positive, even in the limit, may seem counter intuitive, 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 (4) approach one. Proposition 1 thus characterizes equilibrium bias in this limit by equating the *rates* at which the informativeness of the two stages tend to zero as h gets small. Since L is a strictly increasing function, $L(0) = 0$, and the LHS of (5) is positive, the limiting value of bias must be positive. More intuitively, note the trivial fact that, when bias is zero, achieving a second-stage draw against a bias is equally likely as achieving a draw with bias in one's favor. As this statement holds *independently* of the ratio of heterogeneity to noise, setting bias equal to zero nullifies not only the informativeness of a second-stage draw but also the *rate* at which this informativeness changes with h. In other words, a strictly positive bias emerges 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. As depicted in Figure 1, in the limit, bias is chosen to maximizes not the *level* of selective efficiency — as selective efficiency becomes independent of bias — but the *rate* at which selective efficiency increases with the agents' heterogeneity. In the limit, optimal bias thus maximizes the potential gains to selective efficiency from a marginal increase in agents' heterogeneity and with bias set to zero these gains can 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 (4): 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

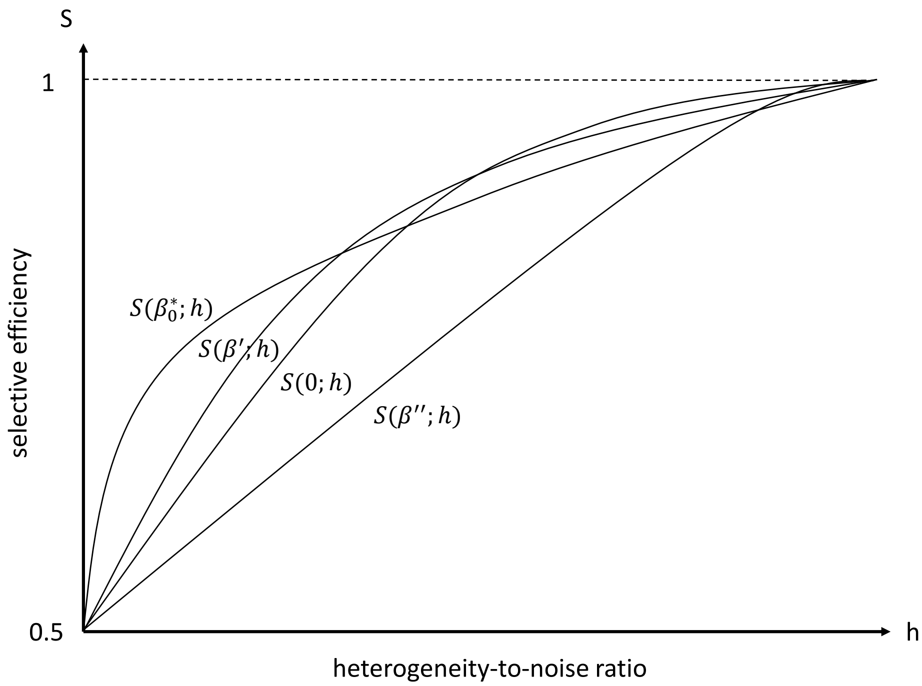

Figure 1: **Selective efficiency.** The figure depicts selective efficiency S as a function of agents' heterogeneity h for different values of bias. $\beta_0^* > 0$ maximizes the slope of S at $h = 0$. When bias is zero or too small ($\beta' < \beta_0^*$) or when bias is too large ($\beta'' >> \beta_0^*$) the potential gains in selective efficiency from a marginal increase in heterogeneity are not fully realized.

illustrated in Figure 2. The left-hand panel plots the density functions for the family of exponential power distributions with mean zero and shape parameter $\alpha > 1$.[13] The right-hand panel in Figure 2 plots the equilibrium bias $\beta^*(h)$ as a function of h, for $\lambda_1 = \lambda_2 = 1$. Despite the myriad possibilities illustrated, we see that, as suggested by Proposition 1, equilibrium bias remains positive even as h gets small for all members of the family.

Our results about the optimal use of bias for selection have implications for our under-

[13] These density functions are given by $g(\Delta\epsilon_t; \alpha) = \frac{\alpha}{2\Gamma(\frac{1}{\alpha})} \exp(-|\Delta\epsilon_t|^\alpha)$, and for all $\alpha > 1$, they satisfy Assumption 1. For $\alpha = 2$, $g(\Delta\epsilon_t; \alpha)$ is a normal distribution with variance $\frac{1}{2}$; as $\alpha \to \infty$, $g(\Delta\epsilon_t; \alpha)$ approaches a uniform distribution with support $[-1, 1]$; and as $\alpha \to 1$, $g(\Delta\epsilon_t; \alpha)$ approaches a Laplace distribution with scale parameter 1. At $\alpha = 1$, Assumption 1 is violated because the Laplace density is not differentiable at 0.

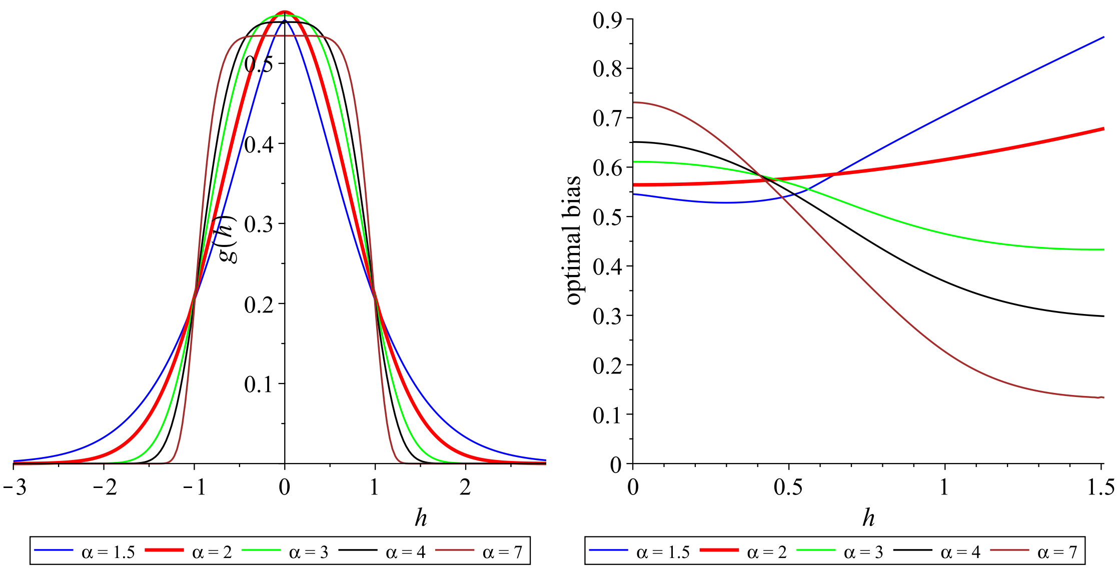

Figure 2: **Example distributions of noise and equilibrium bias.** The left-hand panel depicts the density functions when noise follows an exponential power distributions with mean zero and shape parameter $\alpha \in \{1.5, 2, 3, 4, 7\}$. The right-hand panel plots the corresponding equilibrium bias, as h varies, assuming that the impact of ability on performance is time-invariant, i.e. $\lambda_1 = \lambda_2 = 1$.

standing 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, an important question to ask, is how institutions and organizational practices shape the relation between performance and outcomes. A straight forward but important implication of the introduction of bias is that it makes initial performance have a more persistent effect on final selection. To see this, note that given the principal's equilibrium choice of bias, $\beta^*(h)$, the persistence of the selection process, defined as the probability with which the first-stage winner becomes selected in the second stage, is given by

$$P(\beta^*(h); h) = G(\lambda_1 h)G(\lambda_2 h + \beta^*(h)) + [1 - G(\lambda_1 h)][1 - G(\lambda_2 h - \beta^*(h)]. \tag{6}$$

Certainly, initial performance has a persistent effect on final selection even in the absence of bias, i.e. $P(0; h) > \frac{1}{2}$, because agents' ability-differential, $h > 0$, is time-invariant, making the first-stage winner more likely to also win the second stage. However, in the limit, as $h \to 0$, persistence vanishes unless it is induced "artificially" through the use of bias, i.e.

$$P_0 \equiv \lim_{h \to 0} P(\beta^*(h); h) = G(\beta_0^*), \tag{7}$$

and $P_0 > \frac{1}{2}$, if and only if the limiting value of bias β_0^* is strictly positive. Hence, a direct implication of Proposition 1 is that luck is made persistent. Also note that it follows from (5) and the strict monotonicity of L that, in spite of ability having only negligible impact on performance in the limit, a larger bias is employed and luck is thus made even more persistent when the agents' first-stage performance is relatively more sensitive to ability than second-stage performance, i.e. when $\lambda_1 > \lambda_2$. According to recent evidence, firms show a tendency to allocate to newly hired workers those tasks that are particularly informative about their abilities (Pastorino, 2024). In light of this empirical fact, our results thus indicate that making luck persistent is especially relevant for a firm's "selection of the best" when workers' abilities are judged on the basis of a series of *heterogeneous* tasks.

Finally, because persistence in (6) is increasing in both variables, the fact that, as seen in Figure 2, equilibrium bias $\beta^*(h)$ can be decreasing suggests that, overall, persistence could also be decreasing in h. This can be seen in Figure 3 for the case $\alpha = 7$, where the indirect effect of a decrease in bias is indeed strong enough to overcome the direct positive effect of an increase in heterogeneity, leading to lower levels of persistence for larger values of agents' heterogeneity. This means that the optimal use of bias might induce initial outcomes to have more long-lasting effects in spite of performance being less attributable to agents' abilities. Contrasted with meritocratic principles, these observations are noteworthy and we thus summarize them formally in the following:

Corollary 1 (Persistence) *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* $\frac{\lambda_1}{\lambda_2}$*. Moreover, initial performance can be induced to have greater impact on final selection in situations where performance-differences are less attributable to ability-differentials, i.e. there exist noise distributions* g *and ranges of* h *for which* $P(\beta^*(h); h)$ *is decreasing in* h.

Corollary 1 shows that two apparent violations of meritocratic principles can be rationalized by the very fact that organizations aim to optimize the allocation of resources to the most gifted. Part one shows that making luck persistent, that is, biasing selection in favor of early performers even when initial success is entirely due to luck, turns out to be a necessary feature of selective efficiency. Similarly, part two shows that the use of bias for selection can make final success *less* correlated with initial performance in settings where

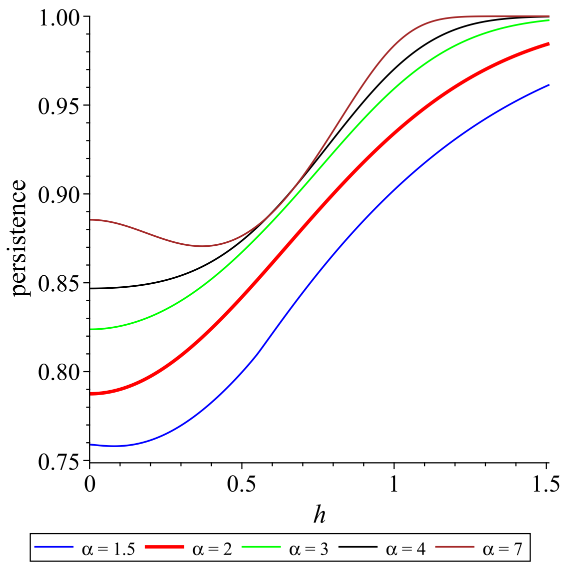

Figure 3: **Equilibrium persistence.** The figure plots the likelihood $P(\beta^*(h); h)$ that winning the initial stage results in becoming ultimately selected, in dependence of the ratio h of agents' heterogeneity to noise. It is assumed that ability has identical impact on performance in all stages, i.e. $\lambda_1 = \lambda_2 = 1$, and that noise follows an exponential power distributions with mean 0 and shape parameter $\alpha \in \{1.5, 2, 3, 4, 7\}$.

performance differentials are *more* attributable to agents' ability-differences. According to our analysis, neither of these features should be considered as an abandonment of meritocratic principles but is, in fact, a direct consequence of the objective of "selecting the best".

4 Strategic behavior of informed agents

Our main model shares with the literature on organizational learning (e.g. Gibbons and Waldman, 2006; Lange, 2007; Pastorino, 2024) the assumption that agents are equally uninformed about their relative abilities as the principle. While the model captures agents' learning, agents learn no faster than the organization. But what if agents have an informational advantage relative to the principal, right from the start? For example, workers might know each other from college or might have shared experiences with previous employers, allowing them to judge their relative abilities. While with uninformed agents, effort choices had no impact on the principal's learning, because, in equilibrium, efforts cancelled each other in the determination of agents' relative performance, when agents are informed, their efforts may no longer be identical. In this section, we determine how informed agents' strategic behavior impacts organizational learning.

Our main concern is that, because effort and ability are *substitutes* in the agents' performance function, the agent thought less likely to be the most able might use the opportunity to make up for his disadvantage by exerting higher effort. If this was the case, small ability differences might be overcome by effort differentials rendering early performance uninformative about agents' abilities not only in the limit when ability differences vanish but also off the limit when ability differences are small. As a consequence, our main result about the relevance of luck for selection would be an artifact of the agents' ignorance of their relative abilities. In this section, we show that, on the contrary, informed agents' strategic behavior does not mitigate but *reinforces* agents' heterogeneity, requiring luck to be made even more persistent than when agents are ignorant of their relative abilities. The following lemma represents the crucial step in our argument:

Lemma 2 (Effort differentials) *Suppose that agents are better informed about their relative abilities than the principal, i.e.* $q^0 > \frac{1}{2}$*. 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 be the one with the higher ability exerts a strictly larger effort than his rival.*

Let us first explain the intuition for Lemma 2 before discussing its implications for the principal's choice of bias and the resulting persistence of luck. Using the same arguments as before, it is easy to see that, in the second stage, agents exert identical efforts. However, unlike in the case where agents are uninformed, the *level* of effort that both agents exert in the second stage and hence their effort costs now depend on the identity of the first-stage winner. To see this, suppose, for simplicity, that $q^0 = 1$, so that agent A is known to be the more able agent with certainty. Remember that the principal is aware of the agents' knowledge but cannot distinguish agent A from agent B. If agent A wins the first stage then bias will reinforce the agents' ability difference, thereby reducing the likelihood of the pivotal realization of noise $\Delta\epsilon_2$ determining second-stage efforts via $C_2'(e_{A,2}^*) = g(h+\beta) = C_2'(e_{B,2}^*)$. Conversely, if agent A loses the first stage, then bias will mitigate the agents' ability difference, leading to an increase in the likelihood of the pivotal realization of $\Delta\epsilon_2$ and second-stage efforts are determined via $C_2'(e_{A,2}^*) = g(h-\beta) = C_2'(e_{B,2}^*)$. Because $g(h+\beta) < g(h-\beta)$ by log-concavity of g, agent A faces lower second-stage effort costs after winning the first stage than after losing, i.e. 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, i.e. agent B has a "cost-saving disincentive". Since the "rewards" of winning the first stage, arising from the increased probability of becoming selected, are the same for both agents, agent A's overall incentive to exert effort in the first stage is thus greater than agent B's.

Lemma 2 shows that, in equilibrium, informed agents' efforts cancel in the final stage of selection but reinforce the agents' ability differential in the initial stage. Let

$$\Delta e_1^*(\beta, h, q^0) \equiv e_{A,1}^*(\beta, h, q^0) - e_{B,1}^*(\beta, h, q^0) \tag{8}$$

denote the agents' equilibrium effort-differential in the first-stage when agents anticipate bias β, have prior q^0 about their relative abilities, and their degree of potential heterogeneity is h. The following result extends Proposition 1 to the case of informed agents under the additional assumption that effort costs are quadratic.

Proposition 2 (Bias with informed agents) *Suppose that $C_t(e_{i,t}) = e_{i,t}^2$ for all i, t, and that agents are better informed about their relative abilities than the principal, i.e. $q^0 > \frac{1}{2}$. In the limit, as noise swamps ability-differences, equilibrium bias $\beta_0^*(q^0) \equiv \lim_{h \to 0} \beta^*(h, q^0)$ is strictly positive, strictly increasing in q^0, and approaches the unique solution to:*

$$2g(0)\left[\lambda_1 + (2q^0 - 1)\frac{\partial \Delta e_1(\beta_0^*(q^0), 0, q^0)}{\partial h}\right] = \lambda_2 L(\beta_0^*(q^0)). \tag{9}$$

Proposition 2 shows that our insights about the optimal use of bias for selection are qualitatively unchanged in the presence of informational asymmetries. When agents are better informed than the principal, the resulting effort-differential makes first-stage performance a more informative signal about agents' relative abilities. As a consequence, the principal will increase the informativeness of the agents' second-stage ranking by employing an even larger bias than in the case where agents are equally uninformed as the principal. Proposition 2 shows that bias remains positive in the limit, even when agents are informed about their relative abilities and can overcome potential ability disadvantages through the exertion of effort. In addition to establishing robustness, the comparative statics contained in Proposition 2 offer further insights about the relevance of luck for selection which we formulate as the following:

Corollary 2 (Persistence amplified) *When informed agents act strategically to maximize their chance to become selected, luck is made even more persistent than when agents are uninformed about their relative abilities or unable to influence their performance through effort, i.e. $P_0(q^0) = G(\beta_0^*(q^0)) > P_0$ for all $q^0 > \frac{1}{2}$. The persistence of luck is amplified by the principal's informational disadvantage, i.e. $P_0(q^0)$ is increasing.*

Corollary 2 emphasizes that making luck persistent can be understood as an organizational response to an informational friction. Organizations employ bias for selection even in extremely noisy environments not only because they know *little* about agents' relative abilities but also because they know *less* than agents themselves. Moreover, because with uninformed agents, i.e. for $q^0 = \frac{1}{2}$, the persistence of luck takes the same value P_0 as in the hypothetical situation where agents cannot influence their performance through effort, Corollary 2 relates our theory to an ongoing discussion of what constitutes "merit" (Sen, 2000). Inherited talents, acquired abilities, and costly noble acts all serve as potential candidates to be combined in this single variable, endowing its possessor with a justification for 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. According to our analysis in Section 3, 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 2 and its corollary 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 "responsability" for their performance.

5 Performance–measurement

Our theory has shown that an organization's "selection of the best" requires luck to be made persistent and that this need becomes more pressing when its agents are informed about their relative abilities and can employ costly efforts to manipulate their chance to become selected. The resulting correlation between the initial and the final outcomes of the dynamic selection process has its origin in the principal's limitation to noisy performance–measurement. As performance–measurement is central to our theory of persistence, in this section we consider how the organization's use of bias varies with the precise way in which performance is measured. For the analysis below and for the remainder of this article, we abstract from the strategic behavior of informed agents considered in Section 4, i.e. we resort to the case where $q^0 = \frac{1}{2}$.

Our model has assumed that performance–measurement is *ordinal* rather than *cardinal*, in that the principal can observe only the *ranking* of the agents' performances, at each stage. Ordinal performance–measurement is prevalent towards the top of an organization's hierarchy, due to the difficulty to quantify performance of increasingly complex tasks (Lazear, 2018). This means that in situations where selection matters most, ordinal performance–measurement is, arguably, the most relevant case to consider. In this section, we consider the counter-factual situation where the principal has access to performance–

information that is *cardinal.* We are especially interested in understanding how a switch from cardinal to ordinal performance–measurement affects an organization's use of bias in the limit where performance becomes approximately random. Given the erosion of cardinal performance–measurement towards the top of an organization's hierarchy, this question is equivalent to asking whether luck can be expected to play a larger or a smaller role for selection in positions with higher ranks.

When performance–measurement is cardinal, the principal can condition her choice of bias on the first-stage *margin of victory* $|\Delta x_1| \equiv |x_{A,1} - x_{B,1}|$. Intuitively, a larger margin of victory is more informative about the agents' relative abilities and thus requires the granting of a larger bias. The equilibrium bias, $\beta^*(h)$, when performance measurement is only ordinal, can be thought of, loosely, as a form of average of the optimal cardinal biases $\beta^c(|\Delta x_1|, h)$ as the margin of victory $|\Delta x_1|$ varies. Proposition 3 below, makes this intuition precise for the limiting case where the heterogeneity-to-noise ratio h tends to zero.

Proposition 3 (Bias with cardinal information) *Let* $q^0 = \frac{1}{2}$*. When performance–information is cardinal rather than ordinal, the following features of equilibrium bias hold in the limit* $h \to 0$ *as noise swamps ability-differences:*

(i) Cardinal bias varies with the first-stage margin of victory and is strictly positive in expectation:

$$\lim_{h \to 0} \mathbb{E}[\beta^c(|\Delta x_1|, h)] > 0. \tag{10}$$

(ii) Ordinal bias equals a form of "average" cardinal bias in the sense that

$$L(\beta_0^*) = \mathbb{E}[L(\beta_0^c(|\Delta x_1|))], \tag{11}$$

where $\beta_0^c(|\Delta x_1|) \equiv \lim_{h\to 0} \beta^c, (|\Delta x_1|, h)$*, and this average becomes exact, i.e.* $\beta_0^* = \mathbb{E}[\beta_0^c(|\Delta x_1|)]$*, when noise is normal.*

Although Proposition 3 derives the parallels between ordinal and cardinal bias for arbitrary $\lambda_1, \lambda_2 > 0$, the properties of cardinal bias become particularly transparent when performance in the two stages is equally sensitive to ability, that is, for $\lambda_1 = \lambda_2$. It is then optimal for the principal to select the agent with the higher *aggregate performance*, $x_{i,1}+x_{i,2}$, and this simple selection rule can be implemented by biasing the agents' second-stage performance in favor of the first-stage winner by exactly $|\Delta x_1|$, the first-stage margin of victory. Hence, in the equilibrium with cardinal information and $\lambda_1 = \lambda_2$, bias is

$$\beta^c(|\Delta x_1|, h) = |\Delta x_1|, \quad \forall |\Delta x_1|, h. \tag{12}$$

For $\lambda_1 = \lambda_2$, it is thus straight forward to understand why, as the heterogeneity-to-noise ratio h goes to zero, the limiting value of bias $\beta_0^c(|\Delta x_1|)$ is positive *on average*. This is because, given (12), expected cardinal bias is just the expected first-stage margin of victory, which, as $h \to 0$, approaches $\mathbb{E}[|\Delta\epsilon_1|] > 0$.

A direct implication of Proposition 3 is that with cardinal performance–information, luck is made persistent *on average*, i.e.

$$P_0^c \equiv \lim_{h\to 0} \mathbb{E}[P(\beta^c(|\Delta x_1|, h), h)] = \lim_{h\to 0} \mathbb{E}[G(\beta^c(|\Delta x_1|, h)] > \frac{1}{2}. \tag{13}$$

A question that arises naturally is: Under which kind of performance–information do organizations induce the greater persistence of luck (on average)? To approach this question it is useful to compare the levels of persistence that ordinal bias induces for the examples depicted in Figure 3 with the case of cardinal performance–measurement. For the case $\lambda_1 = \lambda_2$ depicted in the figure, the expected persistence of luck under cardinal information can be determined easily as follows. Note that the first-stage winner i becomes selected as the agent with the highest aggregate performance, when one of two events occur; (1) when he performs better also in the second stage, i.e. $x_{i,2} > x_{j,2}$; or (2) when his rival j performs better but with a smaller margin of victory, i.e. $x_{j,2} > x_{i,2}$ but $|\Delta x_2| < |\Delta x_1|$. In the limit $h \to 0$, these two events occur with probability $\frac{1}{2}$ and $\frac{1}{4}$, respectively, and the expected persistence of luck under cardinal performance–measurement is thus given by

$$P_0^c = \mathbb{P}(\Delta x_1 + \Delta x_2 \geq 0 | \Delta x_1 \geq 0) = \frac{1}{2} + \frac{1}{4} = \frac{3}{4}. \tag{14}$$

Note from Figure 3 that the persistence of luck under ordinal performance–measurement is larger than $\frac{3}{4}$ for all values of α depicted. For $\alpha = 2$ this is not surprising, because, as shown by Proposition 3, ordinal bias is given by the expectation of cardinal bias when noise is normal and persistence in (13) is a concave function of bias.[14] In fact, combining this insight with equation (11), a sufficient condition for luck to be more persistent under ordinal bias than under cardinal bias is that the function $L(.)$ is convex.[15] We summarize these insights in the following:

Corollary 3 (Persistence with cardinal information) *Suppose that $L(.)$ is convex which holds, for instance, if noise is "sufficiently normal". When the organization ob-*

[14]Because the log-concavity of g implies that g is unimodal and because g is assumed to be symmetric around zero, in the positive domain, g has to be decreasing and hence G has to be concave.

[15]To see that L convex is not necessary note that for the exponential power family of distributions in (13), L is convex if and only if $\alpha \geq 2$ but persistence of luck is larger under ordinal than under cardinal information for all $\alpha > 1.38$. The case where L is convex corresponds to densities $\tilde{g}$ that are more log-concave than normal, i.e. where $\ln \tilde{g}$ is a concave transform of $\ln g$ when g is normal.

serves cardinal performance–information it will induce luck to be less persistent on average than when performance-measurement is only ordinal, i.e. $P_0^c = \mathbb{E}[G(\beta_0^c(|\Delta x_1|)] < G(\beta_0^*) = P_0^*$.

Lazear (2000) argues that due to the difficulty to quantify the performance of managers, firms employ incentive contracts that rely on cardinal information, such as piece rates, less frequently for positions at higher ranks and instead use ordinal comparisons to provide incentives. Corollary 3 shows that when organizations observe cardinal performance–information at low ranks whereas performance-information at high ranks is only ordinal, organizations will induce luck to have more persistent effects on selection for positions towards the top of the hierarchy. Surprisingly luck is induced to play a more important role for those positions where the selection of the best candidate is most critical. In particular, our analysis in this section therefore rationalizes the relevance of luck for the selection of an organization's leadership, where $\Pi(a)$ is, arguably, most sensitive to ability differentials and the selection of the "right" or the "wrong" agent can lead to substantial gains or losses.[16]

6 Discrimination and cumulative disadvantage

Inlcude new section here.

7 Conclusion

When the careers of professional hockey players or CEOs are kick-started by the proximity of their birthday to a cut-off or when hedge funds or venture capitalists 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 play into the hands of recent critics of a meritocratic worldview (e.g. Piketty, 2014; Sandel, 2020), which, in spite of forming the basis of modern democratic societies, is claimed to be a myth, used as a justification for their exorbitant degrees of economic and social inequality. The main contribution of this paper is to show that making initial

[16]For instance, the passing of the UK premiership from Neville Chamberlain to Winston Churchill, rather than to Lord Halifax, during the early years of Word War II brought to an end the British policy of appeasement and has been credited as a major contributor to the Allied victory (Roberts, 2019). Conversely, the infamous decline of Kodak has been attributed to the appointment of CEO Kay Whitmore, who was criticized for lacking the visionary foresight of his rival, Phil Samper, concerning the emergence of digital photography. See "How mediocre managers ruined Kodak" available online at https://www.hrmagazine.co.uk/content/features/how-mediocre-managers-ruined-kodak/.

luck to have a persistent effect on selection is consistent with – if not a necessary feature of – a society aiming to allocate resources and decision-making power to the most able individuals.

Our theory is stylized but capable of identifying a basic mechanism rationalizing the persistence of initial luck as an equilibrium outcome of the strategic interaction between an organization aiming to maximize selective efficiency and a group of heterogeneous agents capable of influencing their likelihood of becoming selected through costly efforts. We have identified the scenarios where the role 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 organizations hierarchy, which means that we have identified luck as a determinant of selection where, arguably, selection matters most.

Appendix

Proof of Lemma 1

Use superscripts W and L 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 $q_{\Delta a}^W(\Delta e_1)$ and $q_{\Delta a}^L(\Delta e_1)$ denote the posterior probabilities that $a_A - a_B = \Delta a$, given Δe_1, and given that A won or lost the first stage, respectively. When there is no risk of confusion, we suppress the dependence of the agents' posteriors on Δe_1.

Step 1. We first show that agents exert identical effort in the second stage and that this holds *independently* of Δe_1 in spite of the influence of first-stage efforts on posteriors. In case W, A's and B's first-order conditions determining second-stage efforts are:

$$C_2'(e_{A,2}^W) = q_h^W g\left(h + \beta + \Delta e_2^W\right) + q_0^W g\left(\beta + \Delta e_2^W\right) + q_{-h}^W g\left(-h + \beta + \Delta e_2^W\right) \quad (15)$$

$$C_2'(e_{B,2}^W) = q_h^W g\left(-h - \beta - \Delta e_2^W\right) + q_0^W g\left(-\beta - \Delta e_2^W\right) + q_{-h}^W g\left(h - \beta - \Delta e_2^W\right). \quad (16)$$

By the symmetry of g, the marginal returns to effort are identical, so $e_{A,2}^{*W} = e_{B,2}^{*W}$. An analogous argument shows $e_{A,2}^{*L} = e_{B,2}^{*L}$.

Step 2. In this step, we determine how the agents identical second-stage efforts vary with a potential first-stage effort-differential. In particular, we show that if $\Delta e_1 > 0$, then $e_{A,2}^{*W} > e_{A,2}^{*L}$. Denote by $q_{\Delta a}^0$ the prior probability that $a_A - a_B = \Delta a \in \{-h, 0, h\}$ and note that because a_A and a_B are identically distributed,

$$q^0_{-h} = q^0_h. \tag{17}$$

From this equality and Assumption 1(i) it immediately follows that, if $\Delta e_1 = 0$, then $e^{*W}_{A,2} = e^{*L}_{A,2}$. To show that $e^{*W}_{A,2} > e^{*L}_{A,2}$ if $\Delta e_1 > 0$, note that given $e^{*W}_{A,2} = e^{*W}_{B,2}$ and $e^{*L}_{A,2} = e^{*L}_{B,2}$, $e^{*W}_{A,2}$ and $e^{*L}_{A,2}$ have to satisfy

$$C_2'(e^{*W}_{A,2}) = q^W_h(\Delta e_1)g(h+\beta) + q^W_0(\Delta e_1)g(\beta) + q^W_{-h}(\Delta e_1)g(-h+\beta) \tag{18}$$

$$C_2'(e^{*L}_{A,2}) = q^L_h(\Delta e_1)g(h-\beta) + q^L_0(\Delta e_1)g(-\beta) + q^L_{-h}(\Delta e_1)g(-h-\beta). \tag{19}$$

Given Assumption 1(i), it follows that

$$\begin{aligned} C_2'(e^{*W}_{A,2}) - C_2'(e^{*L}_{A,2}) &= [q^W_h(\Delta e_1) - q^L_{-h}(\Delta e_1)]g(h+\beta) \\ &+ [q^W_{-h}(\Delta e_1) - q^L_h(\Delta e_1)]g(-h+\beta) + [q^W_0(\Delta e_1) - q^L_0(\Delta e_1)]g(\beta). \end{aligned} \tag{20}$$

To complete Step 2, we show that (20) is strictly positive which, combined with the convexity of C_2, implies that $e^{*W}_{A,2} > e^{*L}_{A,2}$. Assumption 1(ii) implies that, for $\Delta e_1 > 0$,

$$q^W_h(\Delta e_1) - q^L_{-h}(\Delta e_1) < 0 \qquad \text{and} \qquad q^W_{-h}(\Delta e_1) - q^L_h(\Delta e_1) > 0. \tag{21}$$

We now show that for $\Delta e_1 > 0$, $q^W_0(\Delta e_1) - q^L_0(\Delta e_1) > 0$. It follows from Assumption 1(i) and condition (17) that $q^W_0(\Delta e_1) > q^L_0(\Delta e_1)$ if and only if

$$q^W_0(\Delta e_1) = \frac{q^0_0 G(\Delta e_1)}{q^0_h[(G(h+\Delta e_1) + G(-h+\Delta e_1)] + q^0_0 G(\Delta e_1)} \tag{22}$$

$$> \frac{q^0_0 G(-\Delta e_1)}{q^0_h[(G(h-\Delta e_1) + G(-h-\Delta e_1)] + q^0_0 G(-\Delta e_1)} = q^W_0(-\Delta e_1) \tag{23}$$

which is equivalent to

$$2G(\Delta e_1) > G(h+\Delta e_1) + G(-h+\Delta e_1). \tag{24}$$

Assumptions 1(i) and 1(ii) imply (a) strict convexity of G on $[-z, 0]$ and (b) strict concavity of G on $[0, z]$. If $-h + \Delta e_1 \geq 0$, (24) follows from (b). Otherwise (a) implies

$$G(-h+\Delta e_1) < \left(\frac{2\Delta e_1}{h+\Delta e_1}\right) G(0) + \left(\frac{h-\Delta e_1}{h+\Delta e_1}\right) G(-h-\Delta e_1), \tag{25}$$

which in turn implies that

$$G(h+\Delta e_1)+G(-h+\Delta e_1) \quad < \quad \left(\frac{2h}{h+\Delta e_1}\right)G(0)+\left(\frac{2\Delta e_1}{h+\Delta e_1}\right)G(h+\Delta e_1). \quad (26)$$

Using $G(-h-\Delta e_1)+G(h+\Delta e_1)=1=2G(0)$, this last term is smaller than $2G(\Delta e_1)$ by (b). This proves that $q_0^W(\Delta e_1)-q_0^L(\Delta e_1)>0$. Returning to (20), Assumption 1(i) and 1(ii) imply that $g(h+\beta)<g(\beta)$ and $g(h+\beta)<g(-h+\beta)$. Also, the three differences in posteriors in square brackets sum to 0. Hence the inequality $q_0^W(\Delta e_1)-q_0^L(\Delta e_1)>0$, combined with those in (21), implies that (20) is strictly positive.

Step 3. In this final step, we argue that for $\Delta e_1>0$, agent B would have a stronger incentive to exert first-stage effort than agent A which leads to a contradiction, allowing us to conclude that $\Delta e_1^*=0$. Consider the overall utility of agent A:

$$\begin{aligned} & q_h^0\{G\left(h+\Delta e_1\right)\left[G\left(h+\beta+\Delta e_2^{*W}\right)-C_2\left(e_{A,2}^{*W}\right)\right] \qquad (27)\\ + \; & \left[1-G\left(h+\Delta e_1\right)\right]\left[G\left(h-\beta+\Delta e_2^{*L}\right)-C_2\left(e_{A,2}^{*L}\right)\right]\} \\ + \; & q_0^0\{G\left(\Delta e_1\right)\left[G\left(\beta+\Delta e_2^{*W}\right)-C_2\left(e_{A,2}^{*W}\right)\right]+\left[1-G\left(\Delta e_1\right)\right]\left[G\left(-\beta+\Delta e_2^{*L}\right)-C_2\left(e_{A,2}^{*L}\right)\right]\} \\ + \; & q_{-h}^0\{G\left(-h+\Delta e_1\right)\left[G\left(-h+\beta+\Delta e_2^{*W}\right)-C_2\left(e_{A,2}^{*W}\right)\right] \\ + \; & \left[1-G\left(-h+\Delta e_1\right)\right]\left[G\left(-h-\beta+\Delta e_2^{*L}\right)-C_2\left(e_{A,2}^{*L}\right)\right]\}-C_1\left(e_{A,1}\right). \end{aligned}$$

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 $\Delta e_2^{*W}=\Delta e_2^{*L}=0$, Assumption 1(i), and condition (17), the marginal benefit of $e_{A,1}$ simplifies to

$$\begin{aligned} & q_h^0\left[g\left(h+\Delta e_1\right)+g\left(-h+\Delta e_1\right)\right]\left\{G\left(h+\beta\right)-C_2\left(e_{A,2}^{*W}\right)-G\left(h-\beta\right)+C_2\left(e_{A,2}^{*L}\right)\right\} \\ & \qquad +q_0^0 g(\Delta e_1)\{G(\beta)-C_2\left(e_{A,2}^{*W}\right)-G\left(-\beta\right)+C_2\left(e_{A,2}^{*L}\right)\} \qquad (28) \end{aligned}$$

Analogously, for agent B the marginal benefit of $e_{B,1}$ becomes

$$\begin{aligned} & q_h^0\left[g\left(h-\Delta e_1\right)+g\left(-h-\Delta e_1\right)\right]\left\{G\left(h+\beta\right)-C_2\left(e_{B,2}^{*L}\right)-G\left(h-\beta\right)+C_2\left(e_{B,2}^{*W}\right)\right\} \\ & \qquad +q_0^0 g(\Delta e_1)\{G(\beta)-C_2\left(e_{B,2}^{*L}\right)-G\left(-\beta\right)+C_2\left(e_{B,2}^{*W}\right)\}. \qquad (29) \end{aligned}$$

By Assumption 1(i) and Step 1, the difference between (28) and (29) has the sign of $C_2(e_{A,2}^{*L})-C_2(e_{A,2}^{*W})$, which by Step 2 is negative when $e_{A,1}-e_{B,1}>0$. But $e_{A,1}-e_{B,1}>0$ implies $C_1'(e_{A,1})-C_1'(e_{B,1})>0$, so such efforts cannot be optimal for both agents. Analogously, $e_{A,1}<e_{B,1}$ would also yield a contradiction. Hence, equilibrium requires

equal first-stage efforts: $e^*_{A,1} = e^*_{B,1}$. These are unique since with $\Delta e_1 = 0$, (28) and (29) are independent of the common level of e_1. ■

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 (3). 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, $\lim_{h\to 0} S_\beta(\beta, h) = 0\ \forall \beta$. Characterizing $\beta^*_0 \equiv \lim_{h\to 0} \beta^*(h)$ thus requires totally differentiating $S_\beta(\beta^*(h); h)$ with respect to h, setting it equal to 0, and letting $h \to 0$. Total differentiation yields

$$\frac{d}{dh} S_\beta(\beta^*(h); h) = S_{\beta h}(\beta^*(h); h) + S_{\beta\beta}(\beta^*(h); h) \frac{\partial \beta^*(h)}{\partial h}, \tag{30}$$

where $\lim_{h\to 0} S_{\beta\beta}(\beta; h) = 0\ \forall \beta$ (since $\lim_{h\to 0} S_\beta(\beta; h) = 0\ \forall \beta$). Hence, (30) and Assumption 1(i) imply that β^*_0 solves

$$\lim_{h\to 0} S_{\beta h}(\beta^*(h); h) = S_{\beta h}(\beta^*_0, 0) = 2\lambda_1 g(0) g(\beta^*_0) + \lambda_2 g'(\beta^*_0) = 0, \tag{31}$$

which gives (5). 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 Lemma 2** ■ **Proof of Proposition 2** ■

Proof of Proposition 3
To abbreviate notation we let $k = |\Delta x_1| \geq 0$ denote the observed first-stage margin of victory.

Part (i) The principal chooses β to maximize $S^c(\beta, k, h)$, the probability of selecting the more able agent, conditional on agents' abilities being different. Conditional on abilities being different, the probability of margin of victory k is $g(k - \lambda_1 h)$ when the stronger agent wins and $g(k + \lambda_1 h)$ when the weaker agent wins. Hence

$$S^c(\beta, k; h) = g(k - \lambda_1 h)G(\lambda_2 h + \beta) + g(k + \lambda_1 h)G(\lambda_2 h - \beta).$$

The corresponding first-order condition is

$$S^c_\beta(\beta, k; h) = g(k - \lambda_1 h)g(\lambda_2 h + \beta) - g(k + \lambda_1 h)g(\lambda_2 h - \beta) = 0 \tag{32}$$

which, by Assumption 1, uniquely determines the optimal cardinal bias $\beta^c(k, h)$ as a strictly increasing function of k, equal to zero for $k = 0$. Since $\lim_{h\to 0} S^c_\beta(\beta, k; h) = 0 \ \forall \beta, k$, characterizing $\beta^c_0(k) \equiv \lim_{h\to 0} \beta^c(k, h)$ requires totally differentiating $S^c_\beta(\beta^c(k, h), k; h)$ with respect to h, setting it equal to zero, and letting $h \to 0$. Steps paralleling the proof of Proposition 1(i) show that $\beta^c_0(k)$ solves $\lim_{h\to 0} S^c_{\beta h}(\beta, k; h) = 0$, which yields

$$L(\beta^c_0(k)) = \frac{\lambda_1}{\lambda_2} L(k). \tag{33}$$

By Assumption 1, $L(0) = 0$ and $L(k) > 0 \ \forall k > 0$. Hence, $\beta^c_0(k) > 0 \ \forall k > 0$. To compute $\mathbb{E}[\beta^c(k, h)]$, remember from the proof of Lemma 1 that the prior probabilities $q^0_{\Delta a}$ that $a_A - a_B = \Delta a \in \{-h, 0, h\}$ satisfy $q^0_{-h} = q^0_h$. The unconditional density of k on its support $[0, z + \lambda_1 h]$ is thus given by

$$2q^0_h g(k - \lambda_1 h) + 2q^0_h g(k + \lambda_1 h) + 2q^0_0 g(k). \tag{34}$$

Hence

$$\begin{aligned}\mathbb{E}[\beta^c(k, h)] &= 2\int_0^{z+\lambda_1 h} \beta^c(k, h)\left[q^0_h g(k - \lambda_1 h) + q^0_h g(k + \lambda_1 h) + q^0_0 g(k)\right] dk \\ &= 2q^0_h \int_{-\lambda_1 h}^{z} \beta^c(v + \lambda_1 h, h) g(v) dv + 2q^0_h \int_{\lambda_1 h}^{z+2\lambda_1 h} \beta^c(v - \lambda_1 h, h) g(v) dv \\ &\quad + 2q^0_0 \int_0^z \beta^c(k, h) g(k) dk, \end{aligned} \tag{35}$$

Since $2q_h^0 + q_0^0 = 1$,

$$\lim_{h \to 0} \mathbb{E}[\beta^c(k,h)] = 2 \int_0^z \beta_0^c(v) g(v) dv.$$

Since $\beta_0^c(k) > 0 \; \forall k > 0$, $\lim_{h \to 0} \mathbb{E}[\beta^c(k,h)] > 0$.

Part (ii) Given (5) and (33), we need only show that $\mathbb{E}[L(k)] = 2g(0)$. As $h \to 0$, (34) 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). ■

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.

Baker, George, Michael Gibbs, and Bengt Holmstrom, "The Internal Economics of the Firm : Evidence from Personnel Data," *Quarterly Journal of Economics*, 1994, *109* (4), 881–919.

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.

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.

Chiappori, Pierre-Andre, Bernard Salanie, and Julie Valentin, "Early Starters versus Late Beginners," *Journal of Political Economy*, 1999, *107* (4), 731–760.

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).

Denrell, Jerker and Chengwei Liu, "Top performers are not the most impressive when extreme performance indicates unreliability," *Proceedings of the National Academy of Sciences of the United States of America*, 2012, *109* (24), 9331–9336.

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.

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.

Frank, Robert H., *Success and Luck: Good Fortune and the Myth of Meritocracy*, Princeton University Press, 2016.

Gallison, Peter, *Einstein's Clocks, Poincare's Maps*, New York: W.W. Norton, 2003.

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.

Gladwell, Malcom, *Outliers: The Story of Success*, Little, Brown and Company, 2008.

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.

Keuschnigg, Marc, Arnout Van De Rijt, and Thijs Bol, "The plateauing of cognitive ability among top earners," *European Sociological Review*, 2023, *20*, 1–14.

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.

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.

Lluis, Stephanie, "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.

Merton, Robert K., "The matthew effect in science," *Science*, 1968, *159* (3810), 56–62.

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.

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 and Human Capital Acquisition," *Journal of Political Economy*, 2024.

Piketty, Thomas, *Capital in the Twenty-First Century*, Cambridge, MA: The Belknap Press of Harvard University Press, 2014.

Roberts, Andrew, *Leadership in War: Essential Lessons from Those Who Made History*, Penguin Random House, 2019.

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 Samuel Bowles Kenneth Arrow and Steven N. Durlauf, eds., *Meritocracy and Economic Inequality*, Princeton (NJ): Princeton University Press, 2000, pp. 5–16.

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.