UiT The Arctic University of Norway

UiT School of Business and Economics
Working Papers in Economics No. 1/20

Competitive Balance: Information Disclosure and Discrimination in an Asymmetric Contest

Derek J. Clark and Tapas Kundu

UiT School of Business and Economics

Working Paper in Economics Number 1/20

September 2020

Competitive Balance: Information Disclosure and Discrimination in an Asymmetric Contest

Derek J. Clark

School of Business and Economics

UiT - The Arctic University of Norway

NO-9037 Tromsø, Norway

derek.clark@uit.no

and

Tapas Kundu

Oslo Business School

Oslo Metropolitan University

Pilestredet 35

NO-0166 Oslo, Norway

tapas.kundu@oslomet.no

https://doi.org/10.7557/22.5618

ISSN: 2535-6534

Abstract

We study a design problem for an effort-maximizing principal in a two-player contest with two dimensions of asymmetry. Players have different skill levels and an information gap exists, as only one player knows the skill difference. The principal has two policy instruments to redress the lack of competitive balance due to asymmetry; she can commit to an information-revealing mechanism, and she can discriminate one of the players by biasing his effort. We characterize the optimal level of discrimination to maximize aggregate effort, showing how this is in turn inextricably linked to the choice of information revelation. Applications are found in newcomer-incumbent situations in an internal labor market, sales-force management, and research contests.

JEL Codes: D02, D72, D82

Keywords: Asymmetric contest; Information design; Discrimination.

1 Introduction

Competition in social, political and economic spheres is often analyzed as a contest in which resources are sunk in order to win a prize. Numerous applications of these frameworks can be found in the literature relating to conflict and warfare, lobbying, elections, internal and external labor markets or various types of research competition.[1] A common theme in much of the existing work is how the characteristics of the competitors and structure of the contest affect the amount of resources or effort used in the competition, and how a contest designer may attempt to influence this; the most usual assumption is that the designer wishes to maximize the resources expended.[2] We consider a contest between an incumbent and a newcomer for a fixed prize. Such situations are often characterized by the newcomer having better information than both the rival and the contest designer about attributes such as own ability that are relevant to the playing of the contest. The difference in this attribute may be large or small, and it is not certain that the incumbent is the superior contestant. Designing the contest to maximize effort in this situation is not a trivial exercise for the principal since it involves negotiating two dimensions of heterogeneity. First, hidden information makes the return to effort uncertain for the uninformed player, discouraging effort.[3] Second, heterogeneity, as captured by a relative skill disparity, is generally acknowledged to be one factor that limits resource use in contest settings (see Chowdhury et al., 2019). A player with a large relative ability, or a large relative valuation of winning, may intimidate opponents into submitting low efforts, and can hence reduce his own efforts and still win with a large probability.[4] In this paper, we set up a simple model that effectively captures the incumbent-newcomer scenario, and in which the principal has two policy instruments at her disposal. She can commit to a signaling mechanism which may reveal - at least partially - the hidden information; furthermore, she can use a policy which treats one of the players preferentially by biasing positively his effort level in the contest. We demonstrate that there is an interesting interplay between these two policy instruments, and that the optimal level and direction of discrimination is inextricably linked to the choice of information revelation.[5]

Our contribution is twofold. First, we analyze a model in which there is asymmetric

[1]See Konrad et al. (2009) and the references therein.

[2]Other aims are possible. In some contexts achieving a close contest may be the objective (Runkel, 2006), or maximizing participation (Azmat and Möller, 2009), or securing the highest quality winner (Serena, 2017b).

[3]Asymmetry in what the players know about the structure of the contest can generally lead to a low effort level (Wärneryd, 2003).

[4]In a dynamic race setting, Konrad and Kovenock (2009) investigate the discouragement effect that arises as one competitor nears the finishing line, causing opponents to simply give up.

[5]Here, the principal chooses whether and how to disclose information. Denter et al. (2020) consider a situation where the agents themselves can choose to reveal information.

information about the abilities of the contestants and show how the skill differential and the existing discrimination policy affects the incentives of the principal to reveal information; second we allow the principal to choose the optimal level of discrimination, in which she has to take account of how this affects the optimal revelation of information. Our model features uncertainty about the characteristics of the contestants, rather than the value of the contested prize. Specifically, there is a skill differential between the two players in carrying out the contest task, the relative value of which is known only to one player, the newcomer; the skill differential can be large or small, and may positively favor either player. This influences the effective level of contest effort which directly affects the success probabilities. The uninformed incumbent has a prior distribution over two possible values of the relative skill. In seeking to maximize the expected effort from the contest, the designer can commit to a set of signals that are sent after the state is determined. Furthermore, we introduce a discrimination parameter which changes the relative productivity of the contestants' effort in determining the probability of success.

When the discrimination parameter is fixed, the designer can influence effort only through the signaling mechanism. We show that the optimal policy of information disclosure depends upon the (fixed) discrimination parameter. When the informed newcomer is discriminated against, it is always optimal to choose a system of signals such that the skill differential is fully disclosed. This is because the discrimination, although it reduces the effect of the skill differential when the uninformed incumbent is skill inferior, has little effect if information is not disclosed: the return to the effort of the uninformed player is then uncertain, and hence effort is risky. Information disclosure eliminates this uncertainty, encouraging the uninformed incumbent to fight even if the opponent is very skilled. When the informed newcomer is positively discriminated, then the magnitude of discrimination is important for the optimal policy of information revelation. With mild positive discrimination of the newcomer, no disclosure is optimal, whilst full disclosure is best for large positive discrimination. In the former case, the players are treated almost the same, and the uninformed incumbent will exert effort as if he faces an average opponent; knowing for certain that the opponent is very skilled or little skilled is detrimental to effort since he will either be discouraged or slack off. When the discrimination is heavily in favor of the informed newcomer, the incumbent has low incentives to exert effort if he thinks he is facing an opponent of average ability; this can be alleviated to some extent by fully revealing the skill differential. For intermediate values of the discrimination parameter, partial disclosure can be optimal even for a binary distribution of the skill differential. If the uninformed incumbent thinks it very likely that the opponent is of low skill level, the designer will find it optimal not to reveal information since the chance that the opponent is highly skilled incites effort. For intermediate levels of discrimination in favor of the informed newcomer, and when the incumbent thinks it likely that the opponent is of high ability, partial disclosure is optimal which involves updating of the prior belief but not full revelation of type.

Rather than just assuming that the discrimination parameter is exogenously fixed and outside of the control of the designer, we allow her to choose this parameter in order to achieve the maximal amount of effort possible. In this optimal choice, she must be mindful of the fact that the choice of the magnitude and direction of discrimination affects the optimal policy of information disclosure. We show that the designer will not want to implement a level of discrimination that involves partial disclosure of the hidden information. She chooses optimally between values of the discrimination parameter for which no disclosure or full disclosure is optimal, and we show how this is connected to the prior beliefs of the uninformed incumbent. When it is thought that the informed newcomer is very likely to be skill-inferior then the designer does not benefit from revealing this to the uninformed opponent, and she chooses to discriminate in favor of the informed (but likely low-skilled) player; we show further how the magnitude of the discrimination depends on the skill level. On the other hand, when the uninformed incumbent thinks that it is likely that the opponent will be highly skilled, the designer must alleviate riskiness of effort for this player by revealing the true state; she will also discriminate against the informed player to encourage effort by both.

One can imagine several applications in which a newcomer competes against an incumbent for a prize, and the ability of the incumbent is known by all, but the newcomer has hidden talent.[6] An internal labor market has features in common with our framework in which an insider and an outsider to a firm compete for a position or a promotion. The outsider knows his own skill level, but the insider is uncertain of the quality of the rival. In the context of a tournament model, Chan (1996) analyses the preferential treatment of one type of candidate depending on the unknown skill level of external workers. If they are expected to be highly skilled, the performance of the firm can be improved by giving preferential treatment to internal candidates; if external candidates are of little threat in terms of skill level, then they can be given an advantage in the promotion contest to incentivize the internal candidates. This is a simple mechanism for leveling the playing field. Our model develops this approach by making manipulation of the information structure a policy instrument in the contest design. In this case, the signaling mechanism that we consider can be likened to the use of aptitude tests by firms to glean some information about their skill level.[7] In sales-force

[6] Denter et al. (2020) also consider this framework, analyzing the information disclosure decisions made by the newcomer himself by choosing to show off his talent, or lie low.

[7] Several large firms use aptitude tests in hiring such as Apple, Samsung, Microsoft and Nike. See https://www.aptitude-test.com/blog/articles/10-major-companies-that-use-aptitude-testing/.

management, an established seller may face competition from an outside challenger and the firm may choose to disclose information about past sales performance of the challenger. Discrimination in this case can be thought of as less administrative duties, better access to back-office resources, more training, and better territories; see, e.g., Skiera and Albers (1998), Farrell and Hakstian (2001), and Krishnamoorthy et al. (2005).

Research contests between competing teams also fit our framework. The effort made by each applicant can be thought of as building up the quality of the team that is competing for a research grant; one team may already be well established whereas a challenger is up-and-coming and not tested. In this case, the research sponsor could grant the new team a pre-project to gather information about its relative skill level.[8] Discrimination in this case could involve preferential treatment of young researchers or aid in writing a good proposal. Many large corporations run internal innovation competitions which see employees (or teams of employees) compete with each other in order to achieve further funding for their projects.[9] Rathi (2014) documents that Thompson Reuters, the US Department of Health and Human Services, Reed Elsevier and TC Transcontinental (the largest printing company in Canada) use different forms of innovation contest, releasing information about previous contests and findings, and using internal mentors as a way of giving an advantage in the competition. Similarly, one may think of procurement contracting as fitting our model, in which an incumbent entrepreneur is well known, whereas a challenger has hidden qualities. Instructing the newcomer to develop a prototype, and revealing the results of its testing, is a way in which information can potentially be garnered and released. Further mechanisms of information revelation may be legally specified. Zhang and Zhou (2016) note that political candidates in the US are required by the Federal Election Campaign Act to reveal campaign contributions and expenditures; this provides information about the financial support base of the candidates.

Related Literature

We draw together three strands of literature in this paper. One relates to discrimination in contests, the second to asymmetric information structures in contests, and the third to the use of signaling mechanisms to reveal hidden information, otherwise known as Bayesian

[8]Serena (2017a) mentions several research contests in which there is an initial stage in which information on the rivals may be gathered and revealed, before final proposals are made. These include the Horizon 2020 submissions to the European Research Council and design competitions run by the Royal Australian Institute of Architects.

[9]See Adamczyk et al. (2012) for a review of research on innovation contests, and Höber (2017) for internal contests.

persuasion (Kamenica and Gentzkow, 2011).[10] Chowdhury et al. (2019) and Mealem and Nitzan (2016) discuss different forms of discrimination (or affirmative action) that are aimed at leveling the playing field to achieve competitive balance in asymmetric contests. Instruments at the disposal of an effort-maximizing principal include among others exclusion of strong players (Baye et al., 1993), caps on efforts (Che and Gale, 1997) or various forms of discrimination such as differential taxation of the prize, or giving head starts or handicaps to some players (see the survey by Mealem and Nitzan, 2016). In the latter case, one may affect the structure of the contest environment through the probability of success by giving a head start or biasing the efforts of one or more players. For the two-player case, Franke (2012) shows how biasing the efforts optimally in relation to competitors' cost of effort (or equivalently prize valuation) leads to maximal contest effort.[11] The optimal bias in this case ensures that both players have the same probability of winning the prize in equilibrium. When there are more than two players, it is likely that weaker ones will prefer not to enter the contest, and this has an effect on the total effort garnered by the designer. Franke et al. (2013) show in this case how a combination of bias affecting effort productivities and head starts can be used to obtain maximal effort. Giving some weaker contestants an incentive to compete means that the probabilities of winning are not equalized in equilibrium, so that the playing field is not perfectly leveled even in the best case scenario for the principal. With more than two players, favoring a player also affects the strategic interaction among other players and the composite effect can be complex. Fu and Wu (2020) consider the design problem for a principal in an n-player lottery contest with heterogeneous contestants differing in their prize valuations. The principal pursues a broad range of objectives that include among others total effort maximization and has two policy instruments $-$ head starts and biases affecting effort productivities. Fu and Wu (2020) show that the multiplicative biases outperform head starts. Furthermore, in the contest designed to obtain maximal effort, the contestants' winning probabilities can be non-monotone with respect to the rankings of their prize valuations. Unlike Franke et al. (2013) and Fu and Wu (2020), we limit our attention to two-player contest under asymmetric information and allow the principal to choose both the extent of information disclosure and the bias affecting effort productivity.

Leaving the framework of complete information, Hurley and Shogren (1998a), Hurley and Shogren (1998b) and Wärneryd (2003) consider how asymmetries in information can affect contest behavior. Particularly relevant to our work is their focus on cases with one-sided informational asymmetry, where one player has better information than the competitor.[12]

[10] Our focus is on a lottery contest, originating in Tullock (1980), rather than an all-pay auction (Hillman and Riley, 1989). See Lu et al. (2018) for an analysis of information disclosure in an all-pay auction.

[11] See also Epstein et al. (2013).

[12] Serena (2017a) in contrast considers a model of information disclosure in a lottery contest with two-sided

Suppose that in a two-player framework, one knows the exact value of the prize but the other knows only the underlying distribution. The effort of the uninformed player is then likened by Hurley and Shogren (1998a) to a risky input which tends to decrease effort in equilibrium, a finding that is reinforced by Wärneryd (2003) who shows that the two-player lottery contest with asymmetric information yields lower equilibrium effort than when the players are symmetrically informed or uninformed. The two-player lottery contest that we investigate incorporates an asymmetry between the players' relative skill at performing the contest task, and an informational asymmetry since the relative skill level is known by only one player.[13] Both dimensions of asymmetry lead to lower efforts in equilibrium, and the task of leveling the playing field in order to encourage players is complex. We consider two policy instruments: one that is directed towards the skill differential, and one that addresses the information gap.

As noted above, in the presence of asymmetric information, the uninformed contestant's belief about the unknown state of skill difference affects both contestants' effort. The designer can improve her payoff by influencing the uninformed contestant's belief. To this end, she commits to a state-conditional distribution of signals before realization of the state. These distributions of signals have the potential to reveal information with varying degrees, and we study the designer's preferences over her information disclosure policy in combination with the discrimination policy. Kamenica and Gentzkow (2011) have operationalized this method of Bayesian persuasion, and it has been applied to a lottery contest by Zhang and Zhou (2016) and Feng and Lu (2016). This latter paper considers information revelation about an unknown number of competitors, whilst Zhang and Zhou (2016) is more relevant for our analysis since it is a two-player contest. The asymmetric information relates to the value of the prize, and the effort-maximizing designer must reveal the state optimally by committing to a signaling mechanism. Kamenica and Gentzkow (2011) show generally that full disclosure is an optimal policy if the payoff of the sender (principal) as a function of the belief of the receiver (uninformed contestant) is globally convex, whilst no disclosure is best when it is globally concave; if the payoff function of the sender has concave and convex portions, then partial disclosure is optimal. Zhang and Zhou (2016) consider first a structure in which the hidden prize value is binary, which yields an expected effort function that is globally convex or concave depending on the valuation by the informed player and the two possible valuations of the uninformed; hence, a signal is optimal that gives either full disclosure of the hidden state, or no disclosure. Only when there are more than two possible valuations can partial disclosure appear, in which the signal reveals the true value

private information.

[13] Brown (2011) notes that both dimensions can be important in determining competitive incentives.

of the prize imperfectly to the uninformed player. Our findings show that even for a binary distribution of the skill differential, partial disclosure can be optimal for some given level of discrimination, although the principal will not implement such discrimination level if she can choose it.[14]

The paper is organized as follows. Section 2 sets up the basic contest and framework for information disclosure, and Section 3 solves for equilibrium effort levels under different informational assumptions. Section 4 considers the optimal information revelation policy for a given level of discrimination, and the optimal direction and magnitude of discrimination is calculated in Section 5. Section 6 concludes. The appendix contains proofs of our results.

2 Model

Two risk-neutral agents, N(newcomer) and I (incumbent), compete for a fixed prize of value 1 in a contest designed by the principal P. The contest score achieved by player I is simply given by his effort e_I, and the score of the newcomer is a multiple of his effort: $\alpha s e_N$, where $\alpha > 0$ and $s > 0$. The parameter α measures the degree of discrimination: $\alpha > 1$ $(0 < \alpha < 1)$ implies that N is positively (negatively) discriminated. This parameter is chosen optimally by the principal in the analysis below. The parameter s is one of the primitives of the model, and measures N's relative skill: $s > 1$ $(0 < s < 1)$ implies that N is superior (inferior) in skill. The success probabilities of N and I, given an effort profile $(e_N, e_I), e_N, e_I > 0$, are given by a player's score relative to the total score in the contest:

$$\rho_N = \frac{\alpha s e_N}{\alpha s e_N + e_I}, \;\; \rho_I = \frac{e_I}{\alpha s e_N + e_I}.$$

This contest success function is commonly used, and has been axiomatized by Clark and Riis (1998).[15] Schaller and Skaperdas (2020) suggest the multiplicative approach taken here as a general way of capturing asymmetry in contests. The skill differential implies that the newcomer is a certain percentage better or inferior than the incumbent at carrying out the contest task, and it would then appear natural that the instrument of the principal should also be multiplicative.[16]

[14] A major difference between our model and Zhang and Zhou (2016) is the fact that our use of the discrimination parameter renders the relative asymmetry between the players a continuous variable, even though the skill differential takes one of two values. For some of these values partial disclosure is optimal. In Zhang and Zhou (2016), it is the relative valuation of the two players that captures the asymmetry, and this takes a finite number of values; either full disclosure or no disclosure is optimal for all of these values in their model.

[15] If $\alpha s e_A + e_B = 0$, we assume that the prize is not awarded. This does not occur in equilibrium, however.

[16] Modeling the skill differential as an additive head start is also a possibility. However, it is widely acknowledged that head starts tend to dampen contest effort in two-player contests (Franke et al., 2013).

We assume that the relative skill s is the source of information asymmetry at the beginning of the game and is referred to as the state of the game. Player N knows the state. Player I and the principal do not know the state, but know that N is fully informed. For simplicity, we assume that there are two possible states, one in which N is superior and the other in which N is inferior. Specifically, s can take only two values in $S := \left\{x, \frac{1}{x}\right\}, x \in (0,1)$, with prior probabilities $q \in (0,1)$ and $(1-q)$, respectively. Since S contains only two values, a distribution over the state space can be expressed with a scaler $p \in [0,1]$ such that $p = \Pr[s = x]$. We will follow this convention, unless stated otherwise. In addition, we use the notation p to denote a generic distribution wherever needed, while q always refers to the prior, which is a parameter of the model.

2.1 Information disclosure

Before the state is realized, the principal commits to and publicly discloses a pair of state-conditional signal distributions $\left\{\pi(\cdot \mid x), \pi\left(\cdot \mid \frac{1}{x}\right)\right\}$ such that $\pi(\cdot \mid s) \in \triangle(M)$, where M is a finite set of signals and $\triangle(M)$ is the set of all probability distributions over M. Once the state s is realized and revealed to N, nature draws a signal $m \in M$ from the distribution $\pi(\cdot \mid s)$. Both agents observe the signal and the uninformed agent I updates his belief. We let q_m denote I's posterior belief that N is inferior, after observing a signal m, i.e., $q_m = \Pr[s = x \mid m]$. Note that since N also observes the signal, N can infer q_m.

The contest then takes places with agents exerting effort simultaneously. The cost of effort is linear and identical for each agent, with a constant marginal cost of one. The agent $i \in \{N, I\}$ chooses effort $e_i \geq 0$ to maximize his expected payoff, denoted by v_i. The principal's payoff is given by the total expected efforts and is denoted by V_P.

The timing of the game is as follows:

1. P chooses the degree of discrimination $\alpha > 0$, which is common knowledge;

2. P commits to and publicly discloses a set of distributions $\left\{\pi(\cdot \mid x), \pi\left(\cdot \mid \frac{1}{x}\right)\right\}$ where $\pi(\cdot \mid s) \in \triangle(M)$, the set of all probability distributions over the signal space M;

3. Nature draws a realization of the state $s \in S$, which is revealed to player N. Next, nature draws a signal $m \in M$ from $\pi(\cdot \mid s)$. The signal m is publicly observed, leading to a posterior belief distribution for I;

With several players they can, however, be a useful tool for encouraging the participation of desirable players (with a high prize valuation for example), and excluding others. With complete information, Franke et al. (2013) show that an effort-maximizing principal can use head starts to limit the set of active players, then employing a multiplicative bias to extract as much surplus as possible from each player. In addition, we show in Section 5 that addressing the asymmetry arising from a multiplicative skill differential with an additive discrimination policy can also be suboptimal.

4. The contest takes place with N and I choosing effort simultaneously.

We study the perfect Bayesian equilibrium of the game.

We make certain assumptions in our model for analytical tractability. First, we consider only two possible states. When there are more than two states, the analysis of optimal information disclosure requires analyzing the convexity property of a function with a multi-dimensional domain. Since we characterize the information disclosure policy for all possible levels of discrimination, this analysis becomes quite intractable. In order to focus on the key issue of the optimal discrimination, we restrict the state space. Further, if both state values have been above (below) one, then the principal can infer that the uninformed player I is always inferior (superior), and the result on the direction of discrimination is straightforward. The more interesting case is the one in which the uninformed player can be either superior or inferior to the informed player. Our assumption that the two state values are reciprocal to each other allows us to analyze this interesting case. The assumption also implies that the relative skill ratio is the same in both states. An undesired consequence of this assumption is that the aggregate effort is independent of the uninformed player's belief if there is no discrimination in favor of any player. The assumption, however, does not limit our analysis in any significant way; this is because discrimination is endogenous in our model and we have demonstrated that the principal never chooses a policy of no discrimination in equilibrium. Our results would be qualitatively unchanged if we consider an alternative specification of the state space: $S = \{x, y\}, 0 < x < 1 < y$. However, by considering $y = 1/x$, we are able to save a parameter and analyze the interplay of skill differential and information disclosure in a tractable manner. Furthermore, note that the principal chooses the discrimination policy at the beginning of the game, so that it is not conditioned on any information that might be revealed about the state. Knowing the possible states of the world, the principal sets the policy in advance, regulating effort through the signaling mechanism; the principal only observes the state if the signal is fully revealing, and we show below that this is not always an optimal choice. We think of the policy choice as being an overarching principle that does not change according to the skill of the player that competes.

3 Contest

We begin at stage 4. At the final stage, the contest can take place under two possible information structures: i) full information and ii) asymmetric information.

3.1 Full information

Suppose both agents know the value of s. Agent i chooses effort to maximize his expected payoff $v_i = \rho_i - e_i$. The following lemma documents the principal's expected payoff in the Nash equilibrium of the game. We use the notation $\mathbb{E}_p [f(s)]$ to denote the expected value of $f(s)$ given that $p = \Pr[s = x]$; therefore, $\mathbb{E}_p [f(s)] := pf(x) + (1-p) f\left(\frac{1}{x}\right)$.

Lemma 1. *In the Nash equilibrium of the full information contest, the principal's ex ante expected payoff is*

$$V_P^F(\alpha, q) = \mathbb{E}_q \left[\frac{2\alpha s}{(1+\alpha s)^2} \right]. \tag{1}$$

3.2 Asymmetric information

Suppose agents are asymmetrically informed. Agent N knows the true value of s. Agent I knows that the opponent is fully informed about the state, but does not know its value himself. Suppose I's belief is given by some $p \in [0,1]$, where $p = \Pr[s = x]$. Agents choose effort to maximize their expected payoffs: $v_N = \rho_N - e_N$, and $v_I = \mathbb{E}_p(\rho_I) - e_I$. The following lemma documents the principal's expected payoff in the Nash equilibrium of the game.

Lemma 2. *In the Nash equilibrium of the asymmetric information contest, the principal's expected payoff is*

$$V_P^A(\alpha, p) = 2\alpha \left(\frac{\mathbb{E}_p\left[\frac{1}{\sqrt{s}}\right]}{\alpha + \mathbb{E}_p\left[\frac{1}{s}\right]} \right)^2. \tag{2}$$

4 Information disclosure with discrimination

4.1 The posterior belief

At stage 3, nature draws a realization of the state $s \in S$ and subsequently, a signal $m \in M$ from $\pi(\cdot \mid s)$, which leads I to update his belief from the prior q to a posterior $q_m \in [0,1]$:

$$q_m = \Pr[s = x \mid m] = \frac{\Pr[m \mid s = x] \Pr[s = x]}{\sum_s \pi(m \mid s) \Pr(s)} = \frac{\pi(m \mid x) q}{\pi(m \mid x) q + \pi\left(m \mid \frac{1}{x}\right)(1-q)}. \tag{3}$$

Observe that a posterior q_m is a random draw from a distribution of $\pi(m \mid s)$ given some state s, which is stochastically drawn from a binary distribution over the state space $S = \left\{x, \frac{1}{x}\right\}$. Therefore, any set of signal distributions $\left\{\pi(\cdot \mid x), \pi\left(\cdot \mid \frac{1}{x}\right)\right\}$ generates a distribution of posteriors $\{q_m\}_{m \in M}$ with probabilities $\sum_{s \in S} \pi(m \mid s) \Pr(s)$. The principal's expected payoff, given the signal distributions $\left\{\pi\left(\cdot \mid x\right), \pi\left(\cdot \mid \frac{1}{x}\right)\right\}$, is

$$\sum_{s\in S}\sum_{m\in M} V_P^A\left(\alpha, q_m\right)\pi\left(m \mid s\right)\Pr\left(s\right). \tag{4}$$

4.2 Optimal information disclosure

At stage 2, the principal's information disclosure policy therefore solves the following problem:

$$\max_{\pi(\cdot|x)\in\triangle(M),\pi\left(\cdot|\frac{1}{x}\right)\in\triangle(M)} \sum_{s\in S}\sum_{m\in M} V_P^A\left(\alpha, q_m\right)\pi\left(m \mid s\right)\Pr\left(s\right) \quad \text{subject to (3).} \tag{5}$$

Following Kamenica and Gentzkow (2011), we reformulate (5) as a constrained optimization problem of choice over posteriors and derive the signal distributions from the optimal posteriors. While this technique is quite general, we illustrate it with an abridged proof, by constructing the signal distributions specific to our context, and it is included in the appendix.

Lemma 3. *The indirect value function of (5) is the same as the indirect value function of the following optimization problem:*

$$\max_{\{q_m\in[0,1],\beta_m\in[0,1]\}_{m\in M}} \sum_{m\in M}\beta_m V_P^A\left(\alpha, q_m\right) \tag{6}$$
$$\textit{subject to} \sum_{m\in M}\beta_m = 1 \textit{ and } \sum_{m\in M}\beta_m q_m = q.$$

Kamenica and Gentzkow (2011) establish that the indirect value function of (6) is well defined and given by the concave closure of the principal's expected payoff under asymmetric information. The concave closure denoted by $Cav\left(\alpha, q\right)$ is the smallest concave function that is everywhere weakly greater than $V_P^A\left(\alpha, q\right)$.[17] Whether or not $V_P^A\left(\alpha, q\right)$ is strictly less than $Cav\left(\alpha, q\right)$ for $q \in (0,1)$, has implication for how information is disclosed. If $V_P^A\left(\alpha, q\right) = Cav\left(\alpha, q\right)$, then the maximum payoff the principal can achieve by committing to some state-conditional signal distributions is exactly the same as what she gets with no information disclosure. If instead $V_P^A\left(\alpha, q\right) < Cav\left(\alpha, q\right)$, then the principal can manipulate I's belief by choosing the signal distributions suitably and improve her expected payoff from $V_P^A\left(\alpha, q\right)$ to $Cav\left(\alpha, q\right)$. The following proposition documents the above findings. The proof

[17] The concave closure is formally defined as follows. Fix α. Let $co\left(V_P^A\left(\alpha, q\right)\right)$ be the convex hull of the graph of $V_P^A\left(\alpha, q\right)$ as a function of q. Then, the concave closure is given by $Cav\left(\alpha, q\right) = \sup\left\{p \mid (p, q) \in co\left(V_P^A\left(\alpha, q\right)\right)\right\}$.

is omitted; see Kamenica and Gentzkow (2011, Proposition 1, Corollaries 1 and 2.) for a general analysis.

Proposition 1 (Kamenica and Gentzkow 2011)**.** *The indirect value of (6), as a function of the prior q, is given by $Cav\,(\alpha, q)$, the concave closure of $V_P^A\,(\alpha, q)$. The principal benefits from adjusting I's belief by disclosing information if and only if $V_P^A\,(\alpha, q) < Cav\,(\alpha, q)$.*

We determine the principal's preferred information disclosure policy from the shape of $V_P^A\,(\alpha, q)$ with respect to q, which is summed up in the following Lemma.

Lemma 4. *Define $\underline{\alpha} := 3 + 2x$ and $\overline{\alpha} := 3 + (2/x)$. The following characterizes the shape of $V_P^A\,(\alpha, q)$:*

(a) *If $0 < \alpha < 1$, then $V_P^A\,(\alpha, q)$ is decreasing and convex in q.*

(b) *If $\alpha = 1$, then $V_P^A\,(\alpha, q)$ is independent of q.*

(c) *If $1 < \alpha \leq \underline{\alpha}$, then $V_P^A\,(\alpha, q)$* is increasing and concave in q.

(d) *If $\underline{\alpha} < \alpha < \overline{\alpha}$, then $V_P^A\,(\alpha, q)$ is increasing and convex in q for $q \in (0, \hat{q})$, and increasing and concave in q for $q \in (\hat{q}, 1)$ where $\hat{q} := (\alpha - \underline{\alpha}) / (\overline{\alpha} - \underline{\alpha})$.*

(e) *If $\overline{\alpha} \leq \alpha$, then $V_P^A\,(\alpha, q)$ is increasing and convex in q.*

Figure 1 plots $V_P^A\,(\alpha, q)$ against q for various values of α. Depending on the curvature of $V_P^A\,(\alpha, q)$, it follows from Lemma 4 that the principal may choose one of the following three information disclosure policy in equilibrium − (i) *Full information disclosure*, (ii) *No information disclosure*, and (iii) *Partial information disclosure*. Below we describe the signal distributions associated with various information disclosure regimes.

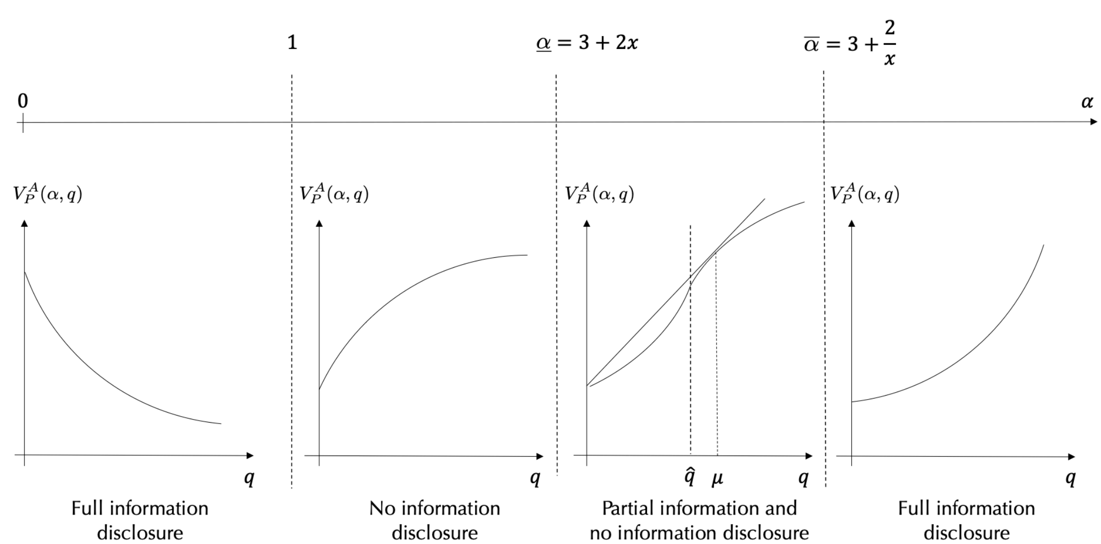

Figure 1: Plot of $V_P^A\,(\alpha, q)$ against q for various values of α

First, consider Lemma 4, cases (**a**) and (**e**), drawn in the first and fourth panels of Figure 1. Since $V_P^A(\alpha, q)$ is convex for all values of $q \in (0,1)$, $Cav(\alpha, q)$ is a straight line joining $V_P^A(\alpha, 0)$ and $V_P^A(\alpha, 1)$ and $V_P^A(\alpha, q) < Cav(\alpha, q)$ except for $q = 0, 1$. The principal benefits from full information disclosure for any $q \in (0,1)$ and the solution of (6) is given by $\beta_1 = q, \beta_2 = 1 - q; q_1 = 1, q_2 = 0$. The principal implements full information disclosure by generating a unique signal in each state so that I perfectly infers the state from observing the signal. Specifically, fix a pair of signals $m_1, m_2 \in M, m_1 \neq m_2$, and consider the following signal distributions, which generate posteriors $q_{m_1} = 1$ and $q_{m_2} = 0$:

$$\pi(m \mid x) = \begin{cases} 1 & \text{if } m = m_1 \\ 0 & \text{if } m \neq m_1 \end{cases} \quad \text{and} \quad \pi\left(m \mid \frac{1}{x}\right) = \begin{cases} 1 & \text{if } m = m_2 \\ 0 & \text{if } m \neq m_2 \end{cases}. \tag{7}$$

Next, consider Lemma 4, case (**c**), drawn in the second panel of Figure 1. Since $V_P^A(\alpha, q)$ is concave for all values of $q \in (0,1)$, $V_P^A(\alpha, q) = Cav(\alpha, q)$ for all q. No information disclosure is optimal and the solution of (6) is given by $\beta_1 = 1; q_1 = q$. The principal implements no information disclosure by generating one and the same signal in every state so that I finds the signal uninformative. Specifically, fix some $m_1 \in M$ and consider the following signal distributions, which generate posterior $q_{m_1} = q$:

$$\pi(m \mid x) = \begin{cases} 1 & \text{if } m = m_1 \\ 0 & \text{if } m \neq m_1 \end{cases} \quad \text{and} \quad \pi\left(m \mid \frac{1}{x}\right) = \begin{cases} 1 & \text{if } m = m_1 \\ 0 & \text{if } m \neq m_1 \end{cases}. \tag{8}$$

Next, consider Lemma 4, case (**d**), depicted in the third panel of Figure 1. In this case, $V_P^A(\alpha, q)$ is partly convex and partly concave for $q \in (0,1)$. There exists a posterior $\mu \in [\hat{q}, 1]$ such that $V_P^A(\alpha, q) < Cav(\alpha, q)$ for $q \in (0, \mu)$ and $V_P^A(\alpha, q) = Cav(\alpha, q)$ for $q \in (\mu, 1)$.[18]

Definition 1. Formally, for any $\alpha > \underline{\alpha}$, $\mu(\alpha)$ is defined as

$$\mu(\alpha) := \begin{cases} p \in (0,1) : V_P^A(\alpha, p) - V_P^A(\alpha, 0) = p\frac{dV_P^A(\alpha, p)}{dp} & \text{if } V_P^A(\alpha, p) - V_P^A(\alpha, 0) > \frac{dV_P^A(\alpha, 1)}{dp} \\ 1 & \text{if } V_P^A(\alpha, p) - V_P^A(\alpha, 0) \leq \frac{dV_P^A(\alpha, 1)}{dp} \end{cases}.$$

Direct calculation gives

$$\mu(\alpha) = \min\left\{\frac{x(\alpha - \underline{\alpha})(x + \alpha)}{(1 - x^2)(1 + 2x + \alpha)}, 1\right\}. \tag{9}$$

[18] It is possible that $\mu = 1$ so that $V_P^I(\alpha, q)$ is always strictly less than $Cav(\alpha, q)$, except for $q = 0, 1$, so that full information disclosure is optimal.

The principal implements partial information disclosure for $q \in (0, \mu(\alpha))$ and no information disclosure for $q \in (\mu(\alpha), 1)$ and the solution of (6) is given by $\beta_1 = \frac{q}{\mu(\alpha)}, \beta_2 = 1 - \frac{q}{\mu(\alpha)}; q_1 = \mu(\alpha), q_2 = 0$ for $q \in (0, \mu)$ and $\beta_1 = 1; q_1 = q$ for $q \in (\mu(\alpha), 1)$.

The principal implements partial information disclosure for $q \in (0, \mu)$ by generating a common signal in both states, while randomizing with another signal in one of the two states. Specifically, fix a pair of signals $m_1, m_2 \in M, m_1 \neq m_2$, and consider the following signal distributions, which generate posteriors $q_{m_1} = \mu$ and $q_{m_2} = 0$:

$$\pi(m \mid x) = \begin{cases} 1 & \text{if } m = m_1 \\ 0 & \text{if } m \neq m_1 \end{cases} \quad \text{and} \quad \pi\left(m \mid \frac{1}{x}\right) = \begin{cases} \frac{q}{\mu(\alpha)(1-q)} & \text{if } m = m_1 \\ 1 - \frac{q}{\mu(\alpha)(1-q)} & \text{if } m = m_2 \\ 0 & \text{if } m \notin \{m_1, m_2\} \end{cases}. \tag{10}$$

Finally, consider Lemma 4, case (**b**). Since $V_P^A(\alpha, q)$ is independent of q, any information disclosure policy yields the same payoff.

The following proposition characterizes the equilibrium information disclosure policy for a given degree of discrimination α. The proof follows from Lemma 4 and the above discussion.

Proposition 2. *Fix $\alpha > 0$, $x \in (0,1)$, and $q \in (0,1)$. The equilibrium information disclosure policy is characterized as follows:*

1. *Suppose I is positively discriminated, i.e., $\alpha < 1$. Then, the principal implements full information disclosure with the signal distributions (7).*

2. *Suppose N is positively discriminated, i.e., $\alpha > 1$.*

   (a) *If $\alpha \leq \underline{\alpha}$, then the principal implements no information disclosure with the signal distributions (8),*

   (b) *If $\underline{\alpha} < \alpha < \overline{\alpha}$, then the principal implements partial information disclosure with the signal distributions (10) for $q \in (0, \mu(\alpha))$, and implements no information disclosure with the signal distributions (8) for $q \in (\mu(\alpha), 1)$,*

   (c) *If $\overline{\alpha} \leq \alpha$, then the principal implements full information disclosure with the signal distributions (7).*

3. *Suppose no player is positively discriminated, i.e., $\alpha = 1$. Then, the principal's expected payoff is invariant to any information disclosure policy.*

Proposition 2 illustrates the interplay between the information advantage and positive discrimination. First, the information advantage and positive discrimination create parallel effects. If the principal gives a candidate preferential treatment, then it is counterproductive to give his competitor an information advantage.[19] Second, the principal's benefits from having a contestant with an information advantage are regressive with respect to the degree of positive discrimination. The second effect arises as both information advantage and positive discrimination create an imbalance among contestants, which has a dampening effect on their contest efforts. The principal can, however, manipulate these two factors of imbalance to address the adverse effect of the inherent disparity in skill.

When the informed newcomer is discriminated against ($\alpha < 1$), it is always optimal to choose a system of signals such that the skill differential is fully disclosed. This is because the discrimination, although it reduces the effect of the skill differential when the uninformed incumbent is skill inferior, has little effect if information is not disclosed: the return to the effort of the uninformed player is then uncertain, and hence effort is risky. Information disclosure eliminates this uncertainty, encouraging the uninformed incumbent to fight even if the opponent is very skilled. When the informed newcomer is positively discriminated, then the magnitude of discrimination is important for the optimal policy of information revelation. With mild positive discrimination of the newcomer ($1 < \alpha \leq \underline{\alpha}$), no disclosure is optimal, whilst full disclosure is best for large positive discrimination ($\overline{\alpha} \leq \alpha$). In the former case, the players are treated almost the same, and the uninformed incumbent will exert effort as if he faces an average opponent; knowing for certain that the opponent is very skilled or little skilled is detrimental to effort since he will either be discouraged or slack off. When the discrimination is heavily in favor of the informed newcomer, the incumbent has low incentives to exert effort if he thinks he is facing an opponent of average ability; this can be alleviated to some extent by fully revealing the skill differential. For intermediate values of the discrimination parameter ($\underline{\alpha} < \alpha < \overline{\alpha}$), these concerns need to be balanced, and partial disclosure is optimal. If the uninformed incumbent thinks it very likely that the opponent is of low skill level ($q \in (\mu(\alpha), 1)$), the designer will find it optimal not to reveal information since the chance that the opponent is highly skilled incites effort. When the incumbent thinks it likely that the opponent is of high ability ($q \in (0, \mu(\alpha))$), partial disclosure is optimal which involves updating of the prior belief but not full revelation of type.

For a given α, the principal's expected payoff from the optimal information disclosure is $Cav(\alpha, q)$. We let FD, ND, and PD denote the sets of values of α, for which the principal

[19] In fact, when the principal positively discriminates I, she would benefit from I having an information advantage − a possibility that our model with one-sided information advantage rules out.

implements full-, no-, and partial-information disclosure in equilibrium, respectively. Then,

$$Cav(\alpha, q) = \begin{cases} V_P^F(\alpha, q) & \text{if } \alpha \in FD \\ V_P^A(\alpha, q) & \text{if } \alpha \in ND \\ \frac{q}{\mu(\alpha)} V_P^A(\alpha, \mu(\alpha)) + \left(1 - \frac{q}{\mu(\alpha)}\right) V_P^A(\alpha, 0) & \text{if } \alpha \in PD \\ \mathbb{E}_q\left[\frac{2s}{(1+s)^2}\right] & \text{if } \alpha = 1 \end{cases} \tag{11}$$

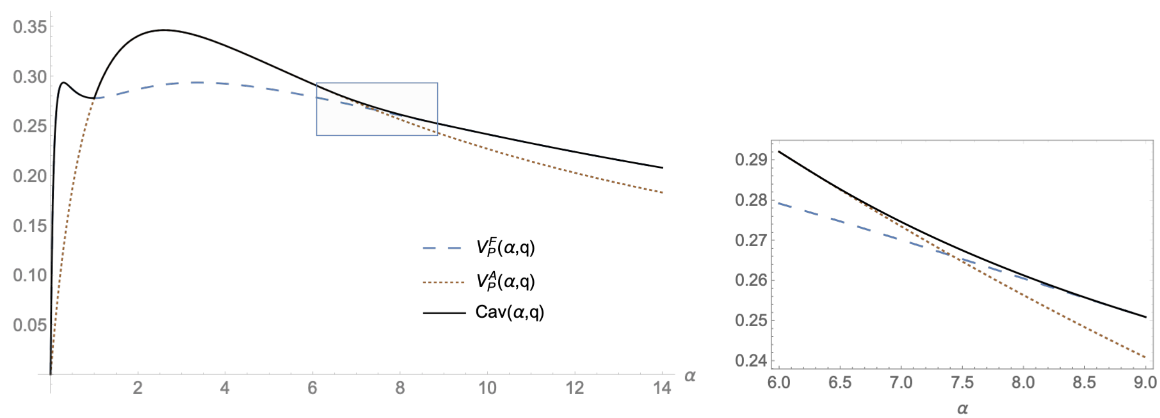

a. $\alpha \in (0, 14]$ *b.* $\alpha \in [6, 9]$

Figure 2: $Cav(\alpha, q)$, $V_P^F(\alpha, q)$, and $V_P^A(\alpha, q)$ against α ($q = 0.5$ and $x = 0.2$)

In Figure 2, panel a and panel b, we plot $Cav(\alpha, q)$, $V_P^F(\alpha, q)$, and $V_P^A(\alpha, q)$ against α, for $q = 0.5$ and $x = 0.2$. In these plots, the continuous curve represents $Cav(\alpha, q)$, the dashed curve represents $V_P^F(\alpha, q)$, and the dotted curve represents $V_P^A(\alpha, q)$. Observe that, in panel a, $Cav(\alpha, q)$ lies weakly above $V_P^F(\alpha, q)$ and $V_P^A(\alpha, q)$. $Cav(\alpha, q)$ can also be strictly higher than $V_P^F(\alpha, q)$ and $V_P^A(\alpha, q)$, in which case, partial information disclosure is optimal. This can be seen in Panel b, which is an expansion of the inset box of Panel a and it considers a sub-range of α, in which $Cav(\alpha, q) > \max\left\{V_P^F(\alpha, q), V_P^A(\alpha, q)\right\}$ for some α.

5 Optimal discrimination

At stage 1, the principal's discrimination policy solves the following problem:

$$\max_{\alpha > 0} \; Cav(\alpha, q). \tag{12}$$

Denote the solution by $\widehat{\alpha}$. We begin with a set of lemmas that will be useful in characterizing $\widehat{\alpha}$.

The first lemma establishes an important link between the discrimination policy and the disparity in skill. Discrimination has two effects. On the one hand, it is costly to the principal due to its asymmetric treatment of the contestants which discourages effort. On the other hand, it benefits the principal by adjusting the disparity in skill, leveling the playing field and encouraging effort. Had there been no uncertainty, the principal could perfectly mitigate the adverse effect of the skill disparity by setting $\alpha = 1/s$. In the presence of uncertainty, discrimination is always costly in one state and can possibly benefit in the other. This fact restricts the magnitude of discrimination employed by the principal, since setting it too high or too low would be costly in both states. The following lemma delineates the optimal discrimination policy, showing that the principal's preferred choice of discrimination is always bounded by the range of the set containing reciprocals of the possible state values, which in our framework is the state space itself.

Lemma 5. *For any $q \in (0,1)$, $\widehat{\alpha} \in \left[x, \frac{1}{x}\right]$.*

A direct implication of Lemma 5 is that the principal will not choose $\alpha > \overline{\alpha}$, a situation described in Lemma 4, case (**e**).

Further, if $\frac{1}{x} \leq \underline{\alpha}$, the principal will also not consider the possibility of partial information disclosure, a situation described in Lemma 4, case (**d**). However, out next result suggests that even if $\underline{\alpha} < \frac{1}{x}$, the principal does not implement partial information disclosure in equilibrium. Specifically, the following lemma shows that the principal's expected payoff in the partial information disclosure regime, given by $\frac{q}{\mu(\alpha)} V_P^A\left(\alpha, \mu\left(\alpha\right)\right) + \left(1 - \frac{q}{\mu(\alpha)}\right) V_P^A\left(\alpha, 0\right)$, is decreasing in α.

Lemma 6. *Consider $\underline{\alpha} < \frac{1}{x}$ and $\alpha \in PD \bigcap \left(\underline{\alpha}, \frac{1}{x}\right]$. Then, the principal's expected payoff from partial information disclosure is decreasing in α.*

From Lemma 6, it immediately follows that if the principal chooses a discrimination policy in $\left(\underline{\alpha}, \frac{1}{x}\right]$ in equilibrium, then it must be associated with no information disclosure; otherwise, the principal can improve her payoff by reducing the discrimination level. The above observation and Proposition 2 together imply that whenever $\widehat{\alpha} \in \left(1, \frac{1}{x}\right]$, the principal implements no information disclosure. This finding is documented in the following corollary. It is worth noting that this result does not imply that the principal's optimal choice of discrimination is always below $\underline{\alpha}$. It can be shown that if, for some q, the principal's expected payoff from no information disclosure reaches its maximum at some $\alpha \in \left(\underline{\alpha}, \frac{1}{x}\right]$, then $\mu\left(\alpha\right) \leq q$. By Proposition 2, it then follows that the principal's payoff from no information disclosure dominates her payoff from partial information disclosure.

Corollary 1. *If $\widehat{\alpha} \in \left(1, \frac{1}{x}\right]$, then the principal implements no information disclosure.*

Whilst in the previous section we demonstrated that partial disclosure can be optimal for some given level of the discrimination parameter, this result shows that the principal will not choose this as part of an optimal policy. She prefers either to remove all uncertainty relating to the state, or to keep the initial level of uncertainty in combination with an optimal discrimination parameter. Simply using the signaling mechanism to change the prior is sub-optimal when the level of discrimination can be chosen.

We denote the principal's optimal choice of discrimination under full and no information disclosure by α^{FD} and α^{ND}, respectively. Formally,

Definition 2. $\alpha^{FD} := \arg\max_{\alpha>0} V_P^F(\alpha, q)$ and $\alpha^{ND} := \arg\max_{\alpha>0} V_P^A(\alpha, q)$.

The following lemmas characterize properties of α^{FD} and α^{ND}.

Lemma 7. *α^{FD} solves $\mathbb{E}_q\left[\frac{s(1-\alpha s)}{(1+\alpha s)^3}\right] = 0$. Further, $\alpha^{FD} \lesseqqgtr 1 \Leftrightarrow q \lesseqqgtr \frac{1}{2}$.*

Lemma 8. *$\alpha^{ND} = \mathbb{E}_q\left[\frac{1}{s}\right]$. Further, $\alpha^{ND} \lesseqqgtr 1 \Leftrightarrow q \lesseqqgtr \frac{x}{1+x}$.*

We now state the main proposition, which describes the equilibrium discrimination policy.

Proposition 3. *Fix $x \in (0, 1)$ and $q \in (0, 1)$. There exists a threshold $\overline{q} \in \left[\frac{x}{1+x}, \frac{1}{2}\right]$ such that*

1. *If $q < \overline{q}$, then the principal's choice of discrimination is $\widehat{\alpha} = \alpha^{FD} < 1$ and there is full information disclosure in equilibrium.*

2. *If $q > \overline{q}$, then the principal's choice of discrimination is $\widehat{\alpha} = \alpha^{ND} > 1$ and there is no information disclosure in equilibrium.*

3. *If $q = \overline{q}$, then the principal is indifferent between choosing $\alpha = \alpha^{FD} < 1$ along with full information disclosure and choosing $\alpha = \alpha^{ND} > 1$ along with no information disclosure.*

Proposition 3 makes several important observations. Consider the principal's preference for information disclosure when she can choose the level of discrimination. Figure 3 plots the equilibrium information disclosure policy in (q, x) space. The asymmetry in information between the players can be both costly and beneficial to the principal. In the presence of asymmetric information, the return to the effort of the uninformed player I is uncertain, which dampens I's incentive to raise effort. By disclosing information, the principal eliminates this uncertainty and encourages I to fight. However, when information is revealed, I comes to know the level of skill disparity in both states. A state-independent choice of discrimination has limited ability to eliminate this disparity; while discrimination may help in leveling the playing field in one state, it makes it more uneven in the other. The disparity in skill discourages I (after becoming informed) to raise effort if he faces a stronger opponent. In such a situation, the principal benefits from not disclosing information.

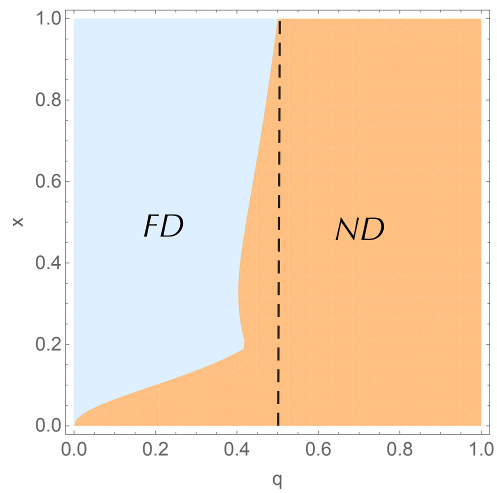

Figure 3: The equilibrium information disclosure policy in (q, x) space

Recall from Corollary 1 that the benefits of not disclosing information outweigh the costs if the informed player N is positively discriminated (i.e., $1 < \alpha < \frac{1}{x}$). The principal can pursue a positive discrimination policy in favor of N to her own advantage if N is more likely to be skill inferior (i.e., $q > 1/2$). Therefore, for $q > 1/2$, the principal commits not to disclose information and positively discriminates N; if I knew that he was facing a low skilled opponent he will slack off, and the principal needs to discriminate N positively to encourage the opponent. If I knew that the opponent is of high skill, then both this and the positive discrimination of N affects contest effort negatively. Hence, the principal does not reveal the state when it is likely that the informed player has low skill. As q decreases further from $1/2$, N becomes more likely to be skill superior. Positive discrimination in favor of N becomes increasingly costly to the principal as I reduces effort anticipating that he is more likely to be facing a stronger and positively discriminated opponent. To counter her disincentive, for sufficiently low values of q, the principal commits to reveal information and positively discriminate I to benefit her in the unfavorable state. However, as is clear from Figure 3, when a low value of q coincides with a low value of x, the policy of no disclosure is optimal; here it is highly likely that the informed player has a very high skill level. Knowing this would dampen I's willingness to exert effort, and then N will also slack-off.

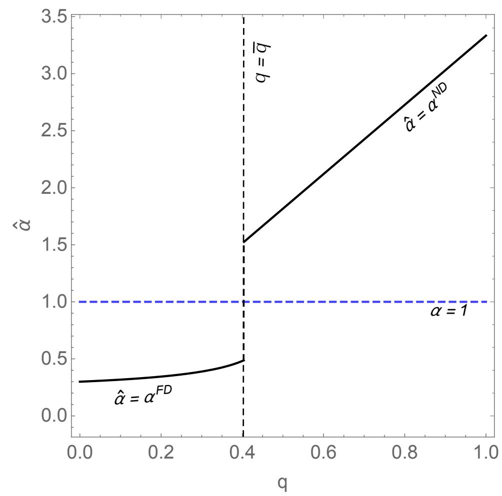

Figure 4: The equilibrium discrimination policy $\widehat{\alpha}$ against q (for $x = 0.3$)

Next, consider the principal's choice of discrimination in equilibrium. Figure 4 plots the equilibrium discrimination $\widehat{\alpha}$ against q, for given $x = 0.3$. If the principal does not disclose information, her optimal choice of discrimination is $\mathbb{E}_q\left[\frac{1}{s}\right] = \alpha^{ND}$. In the Appendix, (A.5) shows that the slope of N's reaction function in equilibrium is positive (negative) when $\alpha > (<)\alpha^{ND}$, so that effort is a strategic complement (substitute) for this player, making him the favorite (underdog).[20] The optimal level of discrimination with no disclosure thus evens the contest, so that the marginal product of N's effort is not affected by I's effort. Similarly, when information is fully disclosed, the principal sets α^{FD}; when the choice of discrimination parameter is made, the state is not known and this level of the discrimination parameter ensures that $\mathbb{E}_q\left[\frac{\partial^2 v_N}{\partial e_I \partial e_N}\right] = 0$, measured in equilibrium.[21] Hence, neither N nor I is the ex ante favorite when α^{FD} is chosen and there is full information disclosure, implying that the contest is expected to be balanced.

In terms of the incumbent-newcomer scenario, we consider that the newcomer (player N) is better informed of relative ability than the incumbent (I). Through previous play, the incumbent is likely to reveal its technological capability, physical and intellectual capacity, resource base or other attributes depending on the application. A newcomer is by definition comparatively unknown. In an innovation contest, the relative capability of a newcomer may

[20]See Hurley and Shogren (1998b). The slope of N's reaction function follows the sign of $\frac{\partial^2 v_N}{\partial e_I \partial e_N}$, the effect that I's effort has on the marginal product of N's own effort. If this is positive then effort is a strategic complement for N.

[21]We have that $\frac{\partial^2 v_N}{\partial e_I \partial e_N} = \alpha s\left[\frac{\alpha s e_N - e_I}{(\alpha s e_N + e_I)^3}\right]$. When the discrimination is set, the state is not known, so taking the expectation of this and evaluating at the symmetric equilibrium gives $sign\,\mathbb{E}_q\left[\frac{\partial^2 v_N}{\partial e_I \partial e_N}\right] = sign\,\mathbb{E}_q\left[\frac{s(\alpha s - 1)}{(\alpha s + 1)^3}\right]$. By Lemma 7, α^{FD} is such that this value is zero.

be determined by undisclosed results of trials or prototype testing, or whether the team is suffering from innovation fatigue or burnout. In a sales contest, the leads that a salesperson has in a certain territory can contribute to relative strength; a worker external to a firm will possess qualities that are initially hidden from a potential employer compared with internal candidates.

In our model, this uncertainty is described by the two parameters (q, x) capturing the likelihood that the newcomer is inferior/superior to the incumbent, and by how much. The principal prefers to fully disclose the hidden information for many parameter combinations where $q < \frac{1}{2}$ so that the newcomer is most likely to be the superior player. Full disclosure can be a requirement to make trial or testing results publicly available in an innovation contest, or to publish territorial leads in a sales contest. Prospective external employees can be subjected to aptitude testing with results that are revealed to all. As we have seen, this policy will be combined with favorable discrimination of the incumbent in the contest. In a sales contest, the incumbent may be given better territories, less administrative duties or other resources in the contest that can boost effective effort. A newcomer in an innovation contest may be required to use administrative resources to prove that its laboratory fulfills certain requirements, provide proof of concept and/or evidence that ethical guidelines are followed. This can reduce effective effort by the newcomer in the contest. An external worker may be required to acquire knowledge of the corporate culture or business practice of the hiring firm, detracting effective effort in the contest. When the incumbent is likely to be inferior (i.e. high values of q), the principal prefers to not reveal information about relative ability, discriminating in favor of the newcomer.

In our model, a multiplicative discrimination policy seems a natural instrument to mitigate the effect of a multiplicative skill differential. However, it can also be shown that a multiplicative discrimination policy is better than an additive discrimination policy from the principal's perspective. To see this, consider the following modification of the game in which the principal chooses an additive discrimination policy: At stage 1 of the game, the principal chooses head starts β_N and β_I so that the respective scores of N and I are given by $se_N + \beta_N$ and $e_I + \beta_I$. Then, the success probabilities of N and I are $\rho_N = \frac{se_N + \beta_N}{se_N + e_I + \beta_N + \beta_I}$ and $\rho_I = \frac{e_I + \beta_I}{se_N + e_I + \beta_N + \beta_I}$. In this modified version of the model, the Nash equilibrium of the asymmetric information contest, given the head starts β_N and β_I, provides the principal with an expected payoff of

$$V_P^A(\beta_N, \beta_I, q) = 2\left(\frac{\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]}{1 + \mathbb{E}_q\left[\frac{1}{s}\right]}\right)^2 - \beta_I - \beta_N \mathbb{E}_q\left[\frac{1}{s}\right].$$

It can be shown that the expected payoff is linear in q. Therefore, the principal's expected payoff is invariant to the information structure and her optimal choices of head starts are $\beta_N = \beta_I = 0$. Further, her maximum expected payoff is $2\left(\frac{\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]}{1+\mathbb{E}_q\left[\frac{1}{s}\right]}\right)^2$, which coincides with her expected payoff from a multiplicative discriminative policy with $\alpha = 1$. Our analysis of the multiplicative discrimination policy shows that the principal's optimal multiplicative discrimination strategy involves either a positive ($\alpha > 1$) or a negative ($\alpha < 1$) discrimination policy in favor of the newcomer and so any additive discrimination strategy is always suboptimal.

6 Concluding remarks

Competition between an incumbent and a newcomer has several defining features which we have captured in a simple model. First, the players have different levels of skills; second they have different information about the skill levels. This affects the contest success function directly, in contrast with previous literature in which the rivals have asymmetric information about the value of the prize. The difference in relative skill level can be large or small, and in favor of the incumbent or the newcomer. Designing the contest to maximize effort is a challenge for the principal, which we have solved by using a combination of two policy instruments. First, she can commit to a signaling mechanism which may reveal - at least partially - the hidden information; second she can use a discrimination policy which treats one of the players preferentially by biasing positively his effort level in the contest. We show that the optimal level and direction of discrimination is linked to the choice of information revelation in a non-trivial way. Further, we show how this is connected to the prior beliefs of the uninformed player. When the uninformed player believes that the informed player is very likely to be skill-inferior, then the designer does not benefit from revealing this to the uninformed opponent, and she chooses to discriminate in favor of the informed (but likely low-skilled) player. On the other hand, when the uninformed player thinks that it is likely that the opponent will be highly skilled, the designer must alleviate riskiness of effort for this player by revealing the true state; she will also discriminate against the informed player to encourage effort by both. The partial disclosure of the hidden information is never optimal when the discriminatory bias can be chosen optimally by the principal. Making the prior more precise is a second-best substitute for optimal discrimination. However, we show that this policy can be optimal for some given levels of the discrimination parameter when this is not under the control of the principal.

Our model captures situations in which a newcomer competes against an incumbent for a prize, and the ability of the incumbent is known by all, but the newcomer has hidden talent. The findings from our analysis can provide insights on how to design information structure to incentivize efforts in these situations. In order to get closed form solutions that allow comparison of effort levels across the whole range of the discrimination parameter, we have made some simplifying assumptions that naturally affect the generality of our results. First, we have assumed that only two players compete; this is, however, an often used construct in contest models. Second, the assumption of a two-point ability distribution is limiting, also due to the fact that the skill levels that can be realized lie either side of the ability of the uniformed player. Hence he knows that he is either facing a superior or an inferior opponent. Nevertheless, we believe that our structure is a useful first step in analyzing the potentially complex interplay between information revelation and other policies designed to level the playing field in contest games.

Appendix

The Appendix contains the proofs.

Proof of Lemma 1:

Proof. Solving the first-order payoff-maximizing conditions of both agents simultaneously. we find that the equilibrium effort under full information is symmetric:

$$e_N^F = e_I^F = \frac{\alpha s}{(1+\alpha s)^2},$$

and the expected payoffs of N and I are:

$$v_N^F = \left[\frac{\alpha s}{1+\alpha s}\right]^2, v_I^F = \left[\frac{1}{1+\alpha s}\right]^2. \tag{A.1}$$

The principal's payoff, expressed as a function of α and q, is given by

$$V_P^F(\alpha, q) = \mathbb{E}_q\left[\frac{2\alpha s}{(1+\alpha s)^2}\right]. \tag{A.2}$$

□

Proof of Lemma 2:

Proof. The optimal effort of N will be a function of the true value of s, whilst the effort of I will be conditioned upon his belief. Equating the marginal benefits of effort for each agent (which equal the common marginal cost of 1) reveals that equilibrium efforts, $e_N^A(s), e_I^A$, satisfy

$$\frac{\alpha s e_I^A}{\left(\alpha s e_N^A(s) + e_I^A\right)^2} = 1 = \mathbb{E}_p\left[\frac{\alpha s e_N^A(s)}{\left(\alpha s e_N^A(s) + e_I^A\right)^2}\right]. \tag{A.3}$$

Then for any value of s,

$$\alpha s e_I^A = \left(\alpha s e_N^A(s) + e_I^A\right)^2,$$

which further implies that

$$\mathbb{E}_p\left[\frac{\alpha s e_N^A(s)}{\alpha s e_I^A}\right] = 1 \Rightarrow e_I^A = \mathbb{E}_p\left[e_N^A(s)\right]. \tag{A.4}$$

The first part of (A.3) implies that

$$e_N^A(s) = \frac{\sqrt{\alpha s e_I^A} - e_I^A}{\alpha s},$$

whereupon taking the expectation gives

$$\mathbb{E}_p\left[e_N^A(s)\right] = \sqrt{e_I^A}\mathbb{E}_p\left[\sqrt{\frac{1}{\alpha s}}\right] - e_I^A\mathbb{E}_p\left[\frac{1}{\alpha s}\right]. \tag{A.5}$$

Using the equality in (A.4) to solve (A.5) gives the final expression for the expected efforts as

$$e_I^A = \alpha\left(\frac{\mathbb{E}_p\left[\frac{1}{\sqrt{s}}\right]}{\alpha + \mathbb{E}_p\left[\frac{1}{s}\right]}\right)^2 = \mathbb{E}_p\left[e_N^A(s)\right].$$

The principal's payoff, expressed as a function of α and p, is given by

$$V_P^A(\alpha, p) = \mathbb{E}_p\left[e_N^A(s)\right] + e_I^A = 2\alpha\left(\frac{\mathbb{E}_p\left[\frac{1}{\sqrt{s}}\right]}{\alpha + \mathbb{E}_p\left[\frac{1}{s}\right]}\right)^2. \tag{A.6}$$

□

Proof of Lemma (3):

Proof. A distribution of posteriors $\{q_m\}_{m\in M}$ with probabilities β_m is *Bayes plausible* if the expected value of the posterior equals the prior, i.e., $\sum_{m\in M}\beta_m q_m = q$. From (3), it follows that any set of state-conditional signal distributions generate a Bayes plausible distribution of posteriors:

$$\sum_{m\in M} q_m \left(\sum_{s\in S} \pi\left(m \mid s\right)\Pr\left(s\right)\right) = \sum_{m\in M} \pi\left(m \mid x\right) q = q. \tag{A.7}$$

The converse is also true − any Bayes plausible distributions of posteriors can be generated from some state-conditional signal distributions. To see this, consider a set of distributions $\{q_m\}_{m\in M}$ with probabilities β_m such that $\sum_{m\in M}\beta_m = 1$ and $\sum_{m\in M}\beta_m q_m = q$. Construct the state-conditional signal distributions as follows: for any given $m \in M$, define

$$\pi\left(m \mid x\right) := \frac{\beta_m q_m}{q} \quad \text{and} \quad \pi\left(m \mid \frac{1}{x}\right) := \frac{\beta_m\left(1-q_m\right)}{1-q}.$$

Then, for any $s \in S$, $\sum_{m\in M}\pi\left(m \mid s\right) = 1$. Further, these constructed conditional signal distributions generate a set of posteriors, which is the same as $\{q_m\}_{m\in M}$:

$$\begin{aligned}\Pr\left[x \mid m\right] &= \frac{\Pr\left[m \mid s=x\right]\Pr\left[s=x\right]}{\Pr\left[m \mid s=x\right]\Pr\left[s=x\right]+\Pr\left[m \mid s=\frac{1}{x}\right]\Pr\left[s=\frac{1}{x}\right]}\\ &= \frac{\frac{\beta_m q_m}{q}\cdot q}{\frac{\beta_m q_m}{q}\cdot q+\frac{\beta_m(1-q_m)}{1-q}\cdot\left(1-q\right)} = q_m.\end{aligned}$$

Therefore, the indirect value function of (5), in which P maximizes her expected payoff over all possible state-conditional signal distributions, is the same as the indirect value function of the following optimization problem, in which P maximizes her expected payoff over all Bayes plausible distribution of posteriors:

$$\begin{aligned}&\max_{\{q_m\in[0,1],\beta_m\in[0,1]\}_{m\in M}} \sum_{m\in M}\beta_m V_P^A\left(\alpha, q_m\right) \\ &\text{subject to} \sum_{m\in M}\beta_m = 1 \text{ and } \sum_{m\in M}\beta_m q_m = q.\end{aligned} \tag{A.8}$$

□

Proof of Lemma 4:

Proof. From (2), we write $V_P^A\left(\alpha,q\right) = 2\alpha\left(\frac{\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]}{\alpha+\mathbb{E}_q\left[\frac{1}{s}\right]}\right)^2$. Consider the first- and the second-order derivatives of $V_P^A(\alpha, q)$ with respect to q:

$$\frac{dV_P^A(\alpha,q)}{dq} = \frac{4\alpha(\alpha-1)\left(\frac{1}{\sqrt{x}}-\sqrt{x}\right)\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]}{\left(\alpha+\mathbb{E}_q\left[\frac{1}{s}\right]\right)^3}, \tag{A.9}$$

$$\frac{d^2V_P^A(\alpha,q)}{dq^2} = \frac{4\alpha(\alpha-1)\left(\frac{1}{\sqrt{x}}-\sqrt{x}\right)^2\left(\alpha-3-2x-2q\left(\frac{1}{x}-x\right)\right)}{\left(\alpha+\mathbb{E}_q\left[\frac{1}{s}\right]\right)^4}. \tag{A.10}$$

From (A.9), it follows that $V_P^A(\alpha, q)$ is increasing, decreasing, and invariant with respect to α if α is greater than, less than, and equal to 1, respectively.

Consider $\alpha < 1$. From (A.10), we find that $\frac{d^2V_P^A(\alpha,q)}{dq^2} > 0$ for all $q \in (0,1)$, implying that $V_P^A(\alpha, q)$ is convex.

Next, consider $\alpha > 1$. From (A.10), we find that $\frac{d^2V_P^A(\alpha,q)}{dq^2} < 0$ for all $q \in (0,1)$ if $\alpha \leq 3+2x$, $\frac{d^2V_P^A(\alpha,q)}{dq^2} > 0$ for all $q \in (0,1)$ if $\alpha \geq 3+\frac{2}{x}$, and $\frac{d^2V_P^A(\alpha,q)}{dq^2} \gtreqless 0$ for all $q \lesseqgtr \frac{\alpha-3-2x}{2\left(\frac{1}{x}-x\right)}$ if $3+2x < \alpha < 3+\frac{2}{x}$. Therefore, $V_P^A(\alpha, q)$ is concave if $\alpha \leq \underline{\alpha} = 3+2x$, convex if $\alpha \geq \overline{\alpha} = 3+\frac{2}{x}$, and convex (concave) in q for $q \in (0, \hat{q})$ (for $q \in (\hat{q}, 1)$) where $\hat{q} = \frac{\alpha-3-2x}{2\left(\frac{1}{x}-x\right)} = \frac{\alpha-\underline{\alpha}}{\overline{\alpha}-\underline{\alpha}}$. □

Proof of Lemma 5:

Proof. Recall that $Cav(\alpha, q)$ is the indirect value function of (A.8). Applying the envelope theorem, we get that

$$\frac{dCav(\alpha,q)}{d\alpha} = \sum_{m\in M}\widehat{\beta}_m\frac{\partial V_P^A(\alpha,\widehat{q}_m)}{\partial\alpha}, \tag{A.11}$$

where $\widehat{\beta}_m \in [0,1]$ and $\widehat{q}_m \in [0,1]$ are solutions of (A.8). Further, for any arbitrary distribution $p \in [0,1]$,

$$\frac{\partial V_P^A(\alpha,p)}{\partial\alpha} = \frac{\partial}{\partial\alpha}2\alpha\left(\frac{\mathbb{E}_p\left[\frac{1}{\sqrt{s}}\right]}{\alpha+\mathbb{E}_p\left[\frac{1}{s}\right]}\right)^2 = 2\mathbb{E}_p^2\left[\frac{1}{\sqrt{s}}\right]\frac{\partial}{\partial\alpha}\left[\frac{\alpha}{\left(\alpha+\mathbb{E}_p\left[\frac{1}{s}\right]\right)^2}\right] = \frac{2\mathbb{E}_p^2\left[\frac{1}{\sqrt{s}}\right]\left[\mathbb{E}_p\left[\frac{1}{s}\right]-\alpha\right]}{\left(\alpha+\mathbb{E}_p\left[\frac{1}{s}\right]\right)^3}, \tag{A.12}$$

which is strictly positive for $\alpha < \mathbb{E}_p\left[\frac{1}{s}\right]$ and strictly negative for $\alpha > \mathbb{E}_p\left[\frac{1}{s}\right]$. Since, for any $\widehat{q}_m \in [0,1]$, $\mathbb{E}_{\widehat{q}_m}\left[\frac{1}{s}\right] \in \left[x, \frac{1}{x}\right]$, (A.11) is strictly positive for $\alpha < x$ and strictly negative for $\alpha > \frac{1}{x}$. Hence, the solution of (12) lies in $\left[x, \frac{1}{x}\right]$. □

Proof of Lemma 6:

Proof. Fix $x \in (0,1)$ and $q \in (0,1)$. Observe that, from Proposition 2, $PD = \{\alpha > \underline{\alpha} : q < \mu(\alpha)\}$. Since $\mu(\alpha) = \min\left\{\frac{x(\alpha-\underline{\alpha})(x+\alpha)}{(1-x^2)(1+2x+\alpha)}, 1\right\}$ is weakly increasing in α, PD is an interval. Further, it can be shown that $\mu\left(\frac{1}{x}\right) < 1$, which implies that for all $\alpha \in PD \bigcap \left(\underline{\alpha}, \frac{1}{x}\right]$, $\mu(\alpha) = \frac{x(\alpha-\underline{\alpha})(x+\alpha)}{(1-x^2)(1+2x+\alpha)}$.

We consider $\alpha \in PD \bigcap \left(\underline{\alpha}, \frac{1}{x}\right]$ and $q \in (0,1)$, and let $V_P^{PD}(\alpha, q)$ denote the principal's expected payoff under partial information disclosure. Further, by Proposition 2, for all $\alpha \in PD$,

$$V_P^{PD}(\alpha, q) = Cav(\alpha, q).$$

Since $Cav(\alpha, q)$ is the indirect value function of (A.8), we apply the envelope theorem to get

$$\frac{dCav(\alpha, q)}{d\alpha} = \sum_{m \in M} \widehat{\beta}_m \frac{\partial V_P^A(\alpha, \widehat{q}_m)}{\partial \alpha},$$

where $\widehat{\beta}_m \in [0,1]$ and $\widehat{q}_m \in [0,1]$ are solutions of (A.8). Note that for all $\alpha \in PD$, these solutions are given by: $\widehat{\beta}_1 = \frac{q}{\mu(\alpha)}, \widehat{\beta}_2 = 1 - \frac{q}{\mu(\alpha)}; \widehat{q}_1 = \mu(\alpha), \widehat{q}_2 = 0$ (see the discussion of Lemma 1, case (**d**), preceding Proposition 2). Therefore, for $\alpha \in PD$, we have

$$\begin{aligned}\frac{dV_P^{PD}(\alpha, q)}{d\alpha} &= \frac{dCav(\alpha, q)}{d\alpha} = \widehat{\beta}_1 \frac{\partial V_P^A(\alpha, \widehat{q}_1)}{\partial \alpha} + \widehat{\beta}_2 \frac{\partial V_P^A(\alpha, \widehat{q}_2)}{\partial \alpha} \\ &= \frac{q}{\mu(\alpha)} \frac{\partial V_P^A(\alpha, \mu(\alpha))}{\partial \alpha} + \left(1 - \frac{q}{\mu(\alpha)}\right) \frac{\partial V_P^A(\alpha, 0)}{\partial \alpha}. \end{aligned} \tag{A.13}$$

We show that both $\frac{\partial V_P^A(\alpha, \mu(\alpha))}{\partial \alpha}$ and $\frac{\partial V_P^A(\alpha, 0)}{\partial \alpha}$ are negative for $\alpha \in PD \bigcap \left(\underline{\alpha}, \frac{1}{x}\right]$.

From (A.12), $\frac{\partial V_P^A(\alpha, 0)}{\partial \alpha} = \frac{2\mathbb{E}_{p=0}^2\left[\frac{1}{\sqrt{s}}\right][x-\alpha]}{\left(\alpha + \mathbb{E}_{p=0}\left[\frac{1}{s}\right]\right)^3} < 0$, since $x < \underline{\alpha} < \alpha$.

Further, from (A.12), $\frac{\partial V_P^A(\alpha, \mu(\alpha))}{\partial \alpha} = \frac{2\mathbb{E}_{\mu(\alpha)}^2\left[\frac{1}{\sqrt{s}}\right]\left[\mathbb{E}_{\mu(\alpha)}\left[\frac{1}{s}\right] - \alpha\right]}{\left(\alpha + \mathbb{E}_{\mu(\alpha)}\left[\frac{1}{s}\right]\right)^3}$, which is negative if and only if $\left[\mathbb{E}_{\mu(\alpha)}\left[\frac{1}{s}\right] - \alpha\right]$ is negative. Note that

$$\begin{aligned}\mathbb{E}_{\mu(\alpha)}\left[\frac{1}{s}\right] - \alpha &= x + \mu(\alpha)\left(\frac{1}{x} - x\right) - \alpha = x + \frac{(\alpha - \underline{\alpha})(x + \alpha)}{(1 + 2x + \alpha)} - \alpha \\ &= \frac{(x + \alpha)(x - \underline{\alpha}) + (1 + x)(x - \alpha)}{(1 + 2x + \alpha)},\end{aligned}$$

which is negative since $x < \underline{\alpha} < \alpha$.

Hence, we conclude that $\frac{dV_P^{PD}(\alpha, q)}{d\alpha}$ is negative for $\alpha \in PD \bigcap \left(\underline{\alpha}, \frac{1}{x}\right]$. ∎

Below we state and prove a result that will be used in the proof of Lemma 7. This result compares the principal's payoffs under full information disclosure from two discrimination policies that are reciprocal to each other.

Lemma A.1. $V_P^F(\alpha,q) \gtreqless V_P^F\left(\frac{1}{\alpha},q\right) \Leftrightarrow \left(q-\frac{1}{2}\right)(\alpha-1) \gtreqless 0.$

Proof of Lemma A.1:

Proof. The following two claims are useful in proving the result.

Claim 1. $V_P^F(\alpha,q) = V_P^F\left(\frac{1}{\alpha},1-q\right)$.

Proof of Claim 1. From (1),

$$V_P^F\left(\frac{1}{\alpha},1-q\right) = \mathbb{E}_{1-q}\left[\frac{2\frac{s}{\alpha}}{\left(1+\frac{s}{\alpha}\right)^2}\right] = \mathbb{E}_{1-q}\left[\frac{2\frac{\alpha}{s}}{\left(\frac{\alpha}{s}+1\right)^2}\right]$$
$$= (1-q)\left[\frac{2\frac{\alpha}{x}}{\left(\frac{\alpha}{x}+1\right)^2}\right] + q\left[\frac{2\alpha x}{(\alpha x+1)^2}\right] = \mathbb{E}_q\left[\frac{2\alpha s}{(1+\alpha s)^2}\right] = V_P^F(\alpha,q).$$

Claim 2. $V_P^F(\alpha,q) \gtreqless V_P^F(\alpha,1-q) \Leftrightarrow (2q-1)(\alpha^2-1) \gtreqless 0.$

Proof of Claim 2: From (1),

$$\begin{aligned}
&V_P^F(\alpha,q) - V_P^F(\alpha,1-q)\\
&= \mathbb{E}_q\left[\frac{2\alpha s}{(1+\alpha s)^2}\right] - \mathbb{E}_{1-q}\left[\frac{2\alpha s}{(1+\alpha s)^2}\right]\\
&= q\left[\frac{2\alpha x}{(\alpha x+1)^2}\right] + (1-q)\left[\frac{2\frac{\alpha}{x}}{\left(\frac{\alpha}{x}+1\right)^2}\right] - (1-q)\left[\frac{2\alpha x}{(\alpha x+1)^2}\right] - q\left[\frac{2\frac{\alpha}{x}}{\left(\frac{\alpha}{x}+1\right)^2}\right]\\
&= (1-2q)\left[\frac{2\frac{\alpha}{x}}{\left(\frac{\alpha}{x}+1\right)^2} - \frac{2\alpha x}{(\alpha x+1)^2}\right] = (1-2q)\left[\frac{2\alpha x}{(\alpha+x)^2} - \frac{2\alpha x}{(\alpha x+1)^2}\right]\\
&= \frac{2\alpha x(1-2q)}{(\alpha+x)^2(\alpha x+1)^2}\left[(\alpha x+1)^2-(\alpha+x)^2\right] = \frac{2\alpha x(1-2q)(1-\alpha^2)(1-x^2)}{(\alpha+x)^2(\alpha x+1)^2}.
\end{aligned}$$

Since $x \in (0,1)$, $V_P^F(\alpha,q) \gtreqless V_P^F(\alpha,1-q) \Leftrightarrow (1-2q)(1-\alpha^2) \gtreqless 0$. This completes the proof of Claim 2.

From Claim 2, replacing α by $\frac{1}{\alpha}$ and q by $1-q$, we get

$$V_P^F\left(\frac{1}{\alpha},1-q\right) \gtreqless V_P^F\left(\frac{1}{\alpha},q\right) \Leftrightarrow (2q-1)\left(\alpha^2-1\right) \gtreqless 0.$$

By Claim 1, it follows then

$$V_P^F(\alpha,q) \gtreqless V_P^F\left(\frac{1}{\alpha},q\right) \Leftrightarrow (2q-1)\left(\alpha^2-1\right) \gtreqless 0 \Leftrightarrow \left(q-\frac{1}{2}\right)(\alpha-1) \gtreqless 0.$$

This completes the proof. $\square$

Proof of Lemma 7:

Proof. Consider the derivative of $V_P^F(\alpha, q)$ with respect to α:

$$\frac{dV_P^F(\alpha,q)}{d\alpha} = \frac{d}{d\alpha}\mathbb{E}_q\left[\frac{2\alpha s}{(1+\alpha s)^2}\right] = \mathbb{E}_q\left[\frac{d}{d\alpha}\frac{2\alpha s}{(1+\alpha s)^2}\right] = \mathbb{E}_q\left[\frac{2s(1-\alpha s)}{(1+\alpha s)^3}\right]. \tag{A.14}$$

Therefore, α^{FD} solves (A.14). By definition of α^{FD}, $V_P^F\left(\alpha^{FD}, q\right) \geq V_P^F\left(\frac{1}{\alpha^{FD}}, q\right)$. It follows from Lemma A.1 that $\left(q-\frac{1}{2}\right)\left(\alpha^{FD}-1\right) \geq 0$. Therefore, $\alpha^{FD} \lesseqgtr 1 \Leftrightarrow q \lesseqgtr \frac{1}{2}$. ☐

Proof of Lemma 8:

Proof. Consider the derivative of $V_P^A(\alpha, q)$ with respect to α:

$$\frac{dV_P^A(\alpha,q)}{d\alpha} = 2\mathbb{E}_q^2\left[\frac{1}{\sqrt{s}}\right]\frac{d}{d\alpha}\left[\frac{\alpha}{\left(\alpha+\mathbb{E}_q\left[\frac{1}{s}\right]\right)^2}\right] = \frac{2\mathbb{E}_q^2\left[\frac{1}{\sqrt{s}}\right]\left[\mathbb{E}_q\left[\frac{1}{s}\right]-\alpha\right]}{\left(\alpha+\mathbb{E}_q\left[\frac{1}{s}\right]\right)^3}. \tag{A.15}$$

Setting the derivative to zero, we get a local optimum at $\alpha = \mathbb{E}_q\left[\frac{1}{s}\right]$. It follows from (A.15) that $V_P^A(\alpha, q)$ is increasing for $\alpha < \mathbb{E}_q\left[\frac{1}{s}\right]$ and decreasing for $\alpha > \mathbb{E}_q\left[\frac{1}{s}\right]$, implying implying that $\mathbb{E}_q\left[\frac{1}{s}\right]$ is a global maximum. Further, $\mathbb{E}_q\left[\frac{1}{s}\right] \lesseqgtr 1 \Leftrightarrow x + q\left(\frac{1}{x}-x\right) \lesseqgtr 1 \Leftrightarrow q \lesseqgtr \frac{x}{1+x}$. ☐

Proof of Proposition 3:

Proof. First, note that Proposition 2, Lemma 5, and Corollary 1 together imply that the principal either chooses $\widehat{\alpha} \in [x, 1]$ and implements full information disclosure, or chooses $\widehat{\alpha} \in \left[1, \frac{1}{x}\right]$ and implements no information disclosure. The following claims, which characterize the principal's optimal choice in various ranges of q, together prove the proposition.

Claim 3. Consider $q > \frac{1}{2}$. Then, $\widehat{\alpha} = \alpha^{ND} > 1$.

Proof of Claim 3: If $\widehat{\alpha} < 1$, by Proposition 2, $Cav(\widehat{\alpha}, q) = V_P^F(\widehat{\alpha}, q)$. Since, $q > \frac{1}{2} \Rightarrow q > \frac{x}{1+x}$, by Lemma 7 and Lemma 8, we have $\alpha^{FD} > 1$ and $\alpha^{ND} > 1$, respectively. Therefore, $\max_{\alpha<1} Cav(\alpha, q) < \max_{\alpha>1} Cav(\alpha, q)$, which proves that $\widehat{\alpha} \not< 1$. Further, $\widehat{\alpha} \neq 1$, since $\alpha^{FD} > 1$ and $\alpha^{ND} > 1$ and $Cav(\alpha, q)$ is continuous at $\alpha = 1$ by the Maximum theorem. Therefore, $\widehat{\alpha} > 1$. It follows from Corollary 1 that $\widehat{\alpha} = \alpha^{ND} > 1$.

Claim 4. Consider $q < \frac{x}{1+x}$. Then, $\widehat{\alpha} = \alpha^{FD} < 1$.

Proof of Claim 4: If $\widehat{\alpha} > 1$, by Corollary 1, $Cav(\widehat{\alpha}, q) = V_P^A(\widehat{\alpha}, q)$. Further, $q < \frac{x}{1+x} \Rightarrow q < \frac{1}{2}$, and therefore, by Lemma 7 and Lemma 8, we have $\alpha^{FD} < 1$ and $\alpha^{ND} < 1$, respectively. Therefore, $\max_{\alpha>1} Cav(\alpha, q) < \max_{\alpha<1} Cav(\alpha, q)$, which proves that $\widehat{\alpha} \not> 1$. Further, $\widehat{\alpha} \neq 1$, since $\alpha^{FD} < 1$ and $\alpha^{ND} < 1$ and $Cav(\alpha, q)$ is continuous at $\alpha = 1$ by the Maximum theorem. Therefore, $\widehat{\alpha} < 1$. It follows from Proposition 2 that $\widehat{\alpha} = \alpha^{FD} < 1$.

Claim 5. For $\frac{x}{1+x} \leq q \leq \frac{1}{2}$, there exists $\overline{q}$ such that $\widehat{\alpha} = \begin{cases} \alpha^{FD} < 1 & \text{if } q < \overline{q} \\ \alpha^{ND} > 1 & \text{if } q > \overline{q} \end{cases}$. Further, the principal is indifferent between choosing $\alpha = \alpha^{FD}$ and choosing $\alpha = \alpha^{ND}$ if $q = \overline{q}$.

Proof of Claim 5: If $\frac{x}{1+x} \leq q \leq \frac{1}{2}$, by Lemma 7 and Lemma 8, we have $\alpha^{FD} \leq 1 \leq \alpha^{ND}$. Since $V_P^F\left(\alpha^{FD}, q\right) = \max_{\alpha>0} V_P^F(\alpha, q)$, we apply the envelope theorem to get

$$\begin{aligned}
\frac{dV_P^F\left(\alpha^{FD}, q\right)}{dq} &= \left.\frac{\partial V_P^F(\alpha, q)}{\partial q}\right|_{\alpha=\alpha^{FD}} = \left.\frac{\partial}{\partial q}\mathbb{E}_q\left[\frac{2\alpha s}{(1+\alpha s)^2}\right]\right|_{\alpha=\alpha^{FD}} \\
&= \frac{2\alpha^{FD}x}{\left(1+\alpha^{FD}x\right)^2} - \frac{2\alpha^{FD}x}{\left(x+\alpha^{FD}\right)^2} \\
&= \frac{2\alpha^{FD}x\left(1-x^2\right)\left(\alpha^{FD}+1\right)\left(\alpha^{FD}-1\right)}{\left(1+\alpha^{FD}x\right)^2\left(x+\alpha^{FD}\right)^2},
\end{aligned}$$

which is negative since $\alpha^{FD} \leq 1$.

Similarly, $V_P^A\left(\alpha^{ND}, q\right) = \max_{\alpha>0} V_P^A(\alpha, q)$, and applying the envelope theorem, we get

$$\begin{aligned}
\frac{dV_P^A\left(\alpha^{ND}, q\right)}{dq} &= \left.\frac{\partial V_P^A(\alpha, q)}{\partial q}\right|_{\alpha=\alpha^{ND}} \\
&= \frac{4\alpha^{ND}\left(\alpha^{ND}-1\right)\left(\frac{1}{\sqrt{x}}-\sqrt{x}\right)\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]}{\left(\alpha^{ND}+\mathbb{E}_q\left[\frac{1}{s}\right]\right)^3} \quad \text{by (A.9)} \\
&= \frac{\mathbb{E}_q\left[\frac{1}{\sqrt{s}}\right]\left(\frac{1}{\sqrt{x}}-\sqrt{x}\right)\left(\alpha^{ND}-1\right)}{2\mathbb{E}_q^2\left[\frac{1}{s}\right]},
\end{aligned}$$

which is positive since $1 \leq \alpha^{ND}$.

The above two observations together imply that $\frac{d}{dq}\left[V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right)\right]$ is positive and therefore $V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right)$ is increasing in q. Since $V_P^F(\alpha, q)$ and $V_P^A(\alpha, q)$ are continuous in q, by the Maximum theorem, $V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right)$ is also continuous in q.

Further, $V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right) > 0$ at $q > \frac{1}{2}$ and $V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right) < 0$ at $q < \frac{x}{1+x}$ by Claims 3 and 4, respectively. Hence, there must exist a threshold $\overline{q} \in \left[\frac{x}{1+x}, \frac{1}{2}\right]$ such that $V_P^A\left(\alpha^{ND}, q\right) - V_P^F\left(\alpha^{FD}, q\right) \lesseqqgtr 0$ if and only if $q \lesseqqgtr \overline{q}$, which proves Claim 5.

Claims 3, 4, and 5 together prove the proposition. $\square$

References

Adamczyk, S., A. C. Bullinger, and K. M. Möslein (2012). Innovation contests: A review, classification and outlook. *Creativity and Innovation Management 21*(4), 335–360. DOI: https://doi.org/10.1111/caim.12003.

Azmat, G. and M. Möller (2009). Competition among contests. *The RAND Journal of Economics 40*(4), 743–768. DOI: https://doi.org/10.1111/j.1756-2171.2009.00087.x.

Baye, M. R., D. Kovenock, and C. G. De Vries (1993). Rigging the lobbying process: An application of the all-pay auction. *American Economic Review 83*(1), 289–294. https://www.jstor.org/stable/2117512.

Brown, J. (2011). Quitters never win: The (adverse) incentive effects of competing with superstars. *Journal of Political Economy 119*(5), 982–1013. DOI: https://doi.org/10.1086/663306.

Chan, W. (1996). External recruitment versus internal promotion. *Journal of Labor Economics 14*(4), 555–570. DOI: https://doi.org/10.1086/209822.

Che, Y.-K. and I. Gale (1997). Rent dissipation when rent seekers are budget constrained. *Public Choice 92*(1-2), 109–126. DOI: https://doi.org/10.1023/A:1017937708549.

Chowdhury, S. M., P. Esteve-González, and A. Mukherjee (2019). Heterogeneity, leveling the playing field, and affirmative action in contests. *unpublished manuscript*.

Clark, D. J. and C. Riis (1998). Contest success functions: an extension. *Economic Theory 11*(1), 201–204. DOI: https://doi.org/10.1007/s001990050184.

Denter, P., J. Morgan, and D. Sisak (2020). Showing off or laying low? the economics of psych-outs. *American Economic Journal: Microeconomics forthcoming.*

Epstein, G. S., Y. Mealem, and S. Nitzan (2013). Lotteries vs. all-pay auctions in fair and biased contests. *Economics & Politics 25*(1), 48–60. DOI: https://doi.org/10.1111/ecpo.12003.

Farrell, S. and A. R. Hakstian (2001). Improving salesforce performance: A meta-analytic investigation of the effectiveness and utility of personnel selection procedures and training interventions. *Psychology & Marketing 18*(3), 281–316. DOI: https://doi.org/10.1002/1520-6793(200103)18:3<281::AID-MAR1009>3.0.CO;2-Z.

Feng, X. and J. Lu (2016). The optimal disclosure policy in contests with stochastic entry: A bayesian persuasion perspective. *Economics Letters 147*, 103–107. DOI: https://doi.org/10.1016/j.econlet.2016.08.038.

Franke, J. (2012). Affirmative action in contest games. *European Journal of Political Economy 28*(1), 105–118. DOI: https://doi.org/10.1016/j.ejpoleco.2011.07.002.

Franke, J., C. Kanzow, W. Leininger, and A. Schwartz (2013). Effort maximization in asymmetric contest games with heterogeneous contestants. *Economic Theory 52*(2), 589–630. DOI: https://doi.org/10.1007/s00199-011-0657-z.

Fu, Q. and Z. Wu (2020). On the optimal design of the biased contest. *Theoretical Economics forthcoming*.

Hillman, A. L. and J. G. Riley (1989). Politically contestable rents and transfers. *Economics & Politics 1*(1), 17–39. DOI: https://doi.org/10.1111/j.1468-0343.1989.tb00003.x.

Höber, B. (2017). *Firm internal innovation contests: Work environment perceptions and employees' participation*. Springer.

Hurley, T. M. and J. F. Shogren (1998a). Asymmetric information contests. *European Journal of Political Economy 14*(4), 645–665. DOI: https://doi.org/10.1016/S0176-2680(98)00030-5.

Hurley, T. M. and J. F. Shogren (1998b). Effort levels in a cournot nash contest with asymmetric information. *Journal of Public Economics 69*(2), 195–210. DOI: https://doi.org/10.1016/S0047-2727(98)00026-7.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615. DOI: https://doi.org/10.1257/aer.101.6.2590.

Konrad, K. A. et al. (2009). *Strategy and Dynamics in Contests*. Oxford, UK: Oxford University Press.

Konrad, K. A. and D. Kovenock (2009). Multi-battle contests. *Games and Economic Behavior 66*, 256–274. DOI: https://doi.org/10.1016/j.geb.2008.05.002.

Krishnamoorthy, A., S. Misra, and A. Prasad (2005). Scheduling sales force training: Theory and evidence. *International Journal of Research in Marketing 22*(4), 427–440. DOI: https://doi.org/10.1016/j.ijresmar.2005.09.001.

Lu, J., H. Ma, and Z. Wang (2018). Ranking disclosure policies in all-pay auctions with incomplete information. *Economic Inquiry 56*(3), 1464–1485. DOI: https://doi.org/10.1111/ecin.12504.

Mealem, Y. and S. Nitzan (2016). Discrimination in contests: a survey. *Review of Economic Design 20*(2), 145–172. DOI: https://doi.org/10.1007/s10058-016-0186-0.

Rathi, A. (2014). To encourage innovation, make it a competition. *Harvard Business Review 11*.

Runkel, M. (2006). Optimal contest design, closeness and the contest success function. *Public Choice 129*(1-2), 217–231. DOI: https://doi.org/10.1007/s11127-006-9030-3.

Schaller, Z. and S. Skaperdas (2020). Bargaining and conflict with up-front investments: How power asymmetries matter. *Journal of Economic Behavior & Organization 176*, 212–225. DOI: https://doi.org/10.1016/j.jebo.2020.01.026.

Serena, M. (2017a). Harnessing beliefs to stimulate efforts. *Working Paper of the Max Planck Institute for Tax Law and Public Finance*. DOI: https://doi.org/10.2139/ssrn.2686543.

Serena, M. (2017b). Quality contests. *European Journal of Political Economy 46*, 15–25. DOI: https://doi.org/10.1016/j.ejpoleco.2016.10.008.

Skiera, B. and S. Albers (1998). Costa: Contribution optimizing sales territory alignment. *Marketing Science 17*(3), 196–213. DOI: https://doi.org/10.1287/mksc.17.3.196.

Tullock, G. (1980). Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock (Eds.), *Toward a theory of the rent-seeking society*, pp. 97–112. College Station: Texas A & M University Press.

Wärneryd, K. (2003). Information in conflicts. *Journal of Economic Theory 110*(1), 121–136. DOI: https://doi.org/10.1016/S0022-0531(03)00006-1.

Zhang, J. and J. Zhou (2016). Information disclosure in contests: A bayesian persuasion approach. *The Economic Journal 126*(597), 2197–2217. DOI: https://doi.org/10.1111/ecoj.12277.