Department of Economics and Finance

Gordon S. Lang School of Business and Economics

University of Guelph

Discussion Paper 2020-03

Microfounded Contest Design

By:

René Kirkegaard
University of Guelph
rkirkega@uoguelph.ca

Department of Economics and Finance
Gordon S. Lang School of Business and Economics | University of Guelph
50 Stone Road East | Guelph, Ontario, Canada | N1G 2W1
www.uoguelph.ca/economics

Microfounded Contest Design*

René Kirkegaard
Department of Economics and Finance
University of Guelph
rkirkega@uoguelph.ca

March 2020

Abstract

This paper examines a unifying model of contests that distinguishes between unobservable actions and observable but noisy performance. Special versions of the model have been used to provide microfoundations for the popular generalized lottery contest success function. However, extensions to contests with exogenous or endogenous biases have strayed from the microfoundations. Consequently, biases and design instruments have been modelled in ad hoc and poorly founded ways. Here, starting directly from the stochastic-performance foundation, internally consistent and fully optimal contests are derived from first principles. The problem resembles a contracting problem. The optimally designed contest is not a generalized lottery contest.

JEL Classification Numbers: C72, D72

Keywords: Contest design; Generalized lottery contests; Stochastic performance contests; Team moral hazard; Tournaments; Tullock contests.

*I thank the Canada Research Chairs programme and SSHRC for funding this research.

1 Introduction

A broad range of economic interactions are contest-like in nature. For the purposes of this paper, think of a contest as an environment in which rival agents take costly actions that influence the probability with which a fixed and indivisible prize is won. Examples include rent-seeking, lobbying, innovation contests, promotion contests, sports, etc. The winner of the prize need not necessarily be the agent who took the most costly action. The mapping from actions to winning probabilities are formalized in the literature by a contest success function (CSF), which may or may not be considered to be a black box.[1]

Generalized lottery or ratio-form CSFs are particularly popular. Here, in an *unbiased contest* with n agents, agent i wins with probability

$$p_i(a_1, a_2, ..., a_n) = \frac{f_i(a_i)}{\sum_{j=1}^n f_j(a_j)}, \tag{1}$$

where $a_j \geq 0$ is agent j's action and where the impact functions $f_j(a_j) \geq 0$ are increasing in a_j. For instance, (1) describes a community raffle where an expenditure of a_i dollars by agent i buys $f_i(a_i)$ raffle tickets. Similarly, Tullock (1980) speaks of a "wealthy eccentric" that for his own reasons sponsors a lottery. Of course, this is just one possible CSF among many.

This paper examines a general and unifying model of contests. Special cases of the model have been explored before. For instance, it is known that the model delivers microfoundations for (1) under additional and restrictive assumptions. It is argued that extensions to *biased contests* have not always stayed true to the premise of these microfoundations. The resulting analysis can be criticized as being ad hoc or poorly founded. This paper provided an internally consistent treatment of the optimal design of biased contests in the general model.

The following simple story is proposed. First, actions are not directly observable but a noisy and observable signal is produced by each agent. Typically, the noisy signal, q_i, can be thought of as the stochastic quality of agent i's performance. In a promotion contest among salespeople, a salesman's performance is

[1] The literature on contests and contest design is enormous. See Konrad (2009), Vojnonić (2015), Mealem and Nitzan (2016), Corchón and Serena (2018), Chowdhury, Esteve-González, and Mukherjee (2019), and Fu and Wu (2019) for recent surveys.

his volume of sales. In innovation contests, a firm's performance is the quality of its innovation. In a competition for a scholarship, a student's performance is his GPA to date. Similarly, a lobbyist's performance is how compelling he can make his agenda or proposal sound. The winner is the agent with the performance of the highest quality. However, the agent's action, a_i, impacts the distribution, $G_i(q_i|a_i)$, of his performance. It is unclear what the most reasonable specification of $G_i(q_i|a_i)$ is, and in any case it is probably sensitive to the application.[2]

The stochastic performance model nests popular CSFs. First, all-pay auctions or deterministic contests trivially arise if $G_i(q_i|a_i)$ is degenerate such that performance and action coincide. Second, in Lazear and Rosen's (1981) rank-order tournament, the action shifts the location of the non-degenerate distribution function. Finally, there are yet other specifications of $G_i(q_i|a_i)$ for which the probability that agent i delivers the best performance reduces to exactly (1).

For instance, (1) materializes if q_i is the best of $f_i(a_i)$ draws from some distribution that is common to all agents, as in Fullerton and McAfee's (1999) research tournament.[3] This and similar microfoundation for (1) are emphasized in e.g. the surveys by Konrad (2009), Vojnonić (2015), Corchón and Serena (2018), and Fu and Wu (2019). However, if this is how a CSF such as (1) is justified, internal consistency demands that any extension that moves beyond unbiased contests must continue to respect the basic stochastic performance premise. This paper seeks to understand the implications of stochastic performance for contest design.

At a very basic level, the premise that actions are unobservable limits the ways in which contests can be manipulated. As explained next, the literature has seemingly ignored this conceptual limitation. However, the assumption that performances are observable provides structure and direction. Thus, it is possible to derive internally consistent and fully optimal contests from first principles.

The standard approach takes the ratio-form in (1) as a jumping-off point and asks how a designer can transform f_i to gainfully manipulate the CSF. It is

[2]Recently, Bastani, Giebe, and Gürtler (2019) have independently proposed a virtually identical model. However, their focus is on comparative statics in unbiased contests.

[3]For other justifications of (1) in this vein, see Hirschleifer and Riley (1992), Clark and Riis (1996), Baye and Hoppe (2003), and Jia (2008). Skapardas (1996) and Clark and Riis (1998) instead take an axiomatic approach to justifying (1). Corchón and Dahm (2011) consider a designer who cannot commit but who is not an expected utility maximizer.

popular to assume that agent i may benefit from an additive bonus, $\delta_i \geq 0$, or a multiplicative bonus, $b_i \geq 0$, or both.[4] Agent i's action is then evaluated by $b_i f_i(a_i) + \delta_i$ rather than just by $f_i(a_i)$. Agent i is said to have a head start over agent j if $\delta_i > \delta_j$ and to be handicapped relative to agent j if $b_i < b_j$. The CSF becomes

$$\widehat{p}_i(a_1, a_2, ..., a_n) = \frac{b_i f_i(a_i) + \delta_i}{\sum_{j=1}^{n} \left(b_j f_j(a_j) + \delta_j\right)}. \tag{2}$$

Implicitly, (2) seems to require that $f_i(a_i)$ is observable. How else can the bonuses be applied to $f_i(a_i)$? It is much less clear where (2) comes from if $f_i(a_i)$ is not observable. Simply put, a compelling microfoundation for (2) is missing.

It is as if (2) is obtained by manipulating (1) directly. However, the CSF is not a primitive of the stochastic performance model. Rather, the CSF is just a reduced form where the uncertainty over performances has been integrated out in order to express winning probabilities as functions only of actions. A sounder approach is to return to the foundation or primitives of the model and build biases into the model from the ground up whenever possible.

Thus, this paper takes direction from the stochastic performance premise to design optimal contests. There are at least two ways of manipulating the contest. Handicapping in golf is a "rules-based" change to the contest; the athlete's performance is observed but then recalculated to determine the winner. Another example is preferential treatment in an admission or promotion contest where the winner need not be the agent with the best entrance score or sales record. These examples require performances to be observable. Handicapping in horse racing as described in Chowdhury, Esteve-González, and Mukherjee (2019) is a "technology-based" intervention where some horses are made to carry extra weight. Here, the unimpeded performance is not observed. This paper focuses on the former type of contest design, which is more easily formalized.[5]

[4] A large literature examines one of these instruments in isolation or both in combination. See e.g. Nti (2004), Epstein, Mealem, and Nitzan (2011), Franke (2012), Franke et. al. (2013), Franke, Leininger, and Wasser (2018), and Fu and Wu (2020). See Mealem and Nitzan (2016) and Chowdhury, Esteve-González, and Mukherjee (2019) for comprehensive surveys of preferential treatment and affirmative action in contests.

[5] Technology-based interventions can often be thought of as manipulating $G_i(q_i|a_i)$. Some contests combine both interventions. First, salesmen may be assigned to different regions or product types, with consequences for $G_i(q_i|a_i)$. Then, sales are compared using some potentially

The idea is to view the problem as a kind of contracting or team moral hazard problem, with the distributions $G_i(q_i|a_i)$ as the primitives. Instead of offering wage schedules as in Holmström (1982), it is winning probabilities that are manipulated to incentivize effort. Thus, the task is to design and commit to an "assignment rule" that maps $(q_1, q_2, ..., q_n)$ into winning probabilities, $P_i(q_1, q_2, ..., q_n)$, subject to incentive compatibility constraints. This is a well-defined and entirely unambiguous problem. Hence, there is no reason a priori to impose ad hoc assumptions on the functional form that the biased CSF must take. Instead, the stochastic performance foundation provides all the structure that is needed to tackle the problem. Likewise, it is not necessary to restrict attention to those $G_i(q_i|a_i)$ that yield (1) in the unbiased case. In sum, the contract theory approach makes it possible to handle stochastic performance in much more generality.

Focus is on a designer whose expected utility depends only on, and is increasing in, the agents' actions. Under familiar technical assumptions, the main result is that for *all* such objective functions, the fundamental structure of the optimal contest is essentially the same. The optimal assignment rule is deterministic and can be implemented by rescaling each agent's likelihood-ratio by an individual-specific factor. The prize is then assigned to the agent with the highest rescaled likelihood-ratio. Thus, the paper identifies a guiding principle for contest design that holds for a large class of distribution functions and objective functions.

Given the optimal assignment rule, the uncertainty over performances can be integrated out to derive the implied endogenous CSF. In Fullerton and McAfee's (1999) model, the resulting CSF does not generically reduce to (2). Hence, a key message is that even if the ratio-form is valid in an unbiased contest, it is not justified to confine attention to the ratio-form when the contest is manipulated.

Although the contribution of the paper is primarily methodological, the way in which contest design is approached has significant implications. Just one example of this comes from the fact that the literature that is based on (2) has concluded that the optimal design in a two-agent contest leads to a completely level playing field. However, this conclusion no longer holds once optimal design is based on the stochastic performance model. The hope is that this paper articulates a way to rigorously approach this and other important questions in contest theory.

---

complicated formula to determine who is promoted. This paper is about the last step.

2 Contests with stochastic performance

This section lays out the basic model. There is a fixed set $N = \{1, ..., n\}$ of contestants or agents. Agent i takes costly action $a_i \in \mathbb{R}_+$. The action influences the distribution of the agent's performance, q_i. The distribution function is written $G_i(q_i|a_i)$. It is assumed to be atomless whenever $a_i > 0$, in which case it has density $g_i(q_i|a_i) > 0$ and support $[\underline{q}_i, \overline{q}_i]$, which may or may not be bounded above or below. Note that the support is the same for all strictly positive actions. If $a_i = 0$, the possibility that the distribution is degenerate at $q_i = \underline{q}_i$ is allowed. Given actions, agents' performances are statistically independent.

In an unbiased contest, the agent with the performance of the highest quality wins. Thus, if $a_i > 0$, agent i's probability of winning is

$$p_i(a_1, ..., a_n) = \int_{\underline{q}_i}^{\overline{q}_i} \left(\prod_{j \neq i} G_j(q|a_j)\right) g_i(q|a_i) dq. \tag{3}$$

Note that (3) at least partly describes a CSF, with the caveat that the winning probability is yet unspecified if $a_i = 0$. If $G_i(q_i|0)$ is non-degenerate, then (3) applies at $a_i = 0$ as well. However, in some applications it may be assumed that $G_i(q_i|0)$ is degenerate such that $q_i = \underline{q}_i$ with probability one. In this case, ties in performances may occur with positive probability. It is assumed that ties are broken with a fair coin. A special case that is of interest is when $G_i(q_i|0)$ is degenerate for all agents and $\underline{q}_i = \underline{q}$ is the same across agents. Then, all agents tie if they all take zero action. In this case, $p_i(0, ..., 0) = \frac{1}{n}$ for all i.

Agent i assigns some exogenous value $v_i > 0$ to winning the contest. The value of losing is zero. Costs are normalized to be linear in a_i. Since costs are increasing in a_i, the action can often be interpreted as effort. Agent i's expected payoff is now

$$U_i(a_1, ..., a_n) = v_i p_i(a_1, ..., a_n) - a_i.$$

Special cases of the model have been considered before. For example, an all-pay auction is a contest with no noise. Here, $G_i(q_i|a_i)$ is degenerate for all actions, such that $q_i = a_i$ with probability one. Then, the agent with the highest action wins. In this paper, such distributions are ruled out. In Lazear and

Rosen's (1981) rank-order tournament, the noise is derived from an agent who can shift the location of the distribution. Formally, $q_i = f_i(a_i) + \varepsilon_i$, where ε_i is the realization of a random variable which is independent of a_i. Thus, unless the support of ε_i is the entire real line, the support of q_i depends on the action a_i. Again, this is ruled out in this paper.

Finally, (3) has been used in the literature to microfound (1). An early contribution in this literature is Hirschleifer and Riley (1992). They propose a model with a multiplicative production function where $q_i = f_i(a_i)\varepsilon_i$, with ε_i exponentially distributed with mean one. With $n = 2$ agents, (3) reduces to (1). The following example allows for more agents.

Example 1 (The best-shot model): Assume that agent i's distribution function can be written

$$G_i(q|a_i) = H_i(q)^{f_i(a_i)},\ q \in [\underline{q}_i, \overline{q}_i], \tag{4}$$

and therefore

$$g_i(q|a_i) = f_i(a_i)H_i(q)^{f_i(a_i)-1}h_i(q),\ q \in [\underline{q}_i, \overline{q}_i],$$

for all $i \in N$, where $H_i(q)$ is a distribution function with density $h_i(q)$. If $f_i(a_i)$ is restricted to take integer values, $G_i(q_i|a_i)$ is the distribution of the best draw from $H_i(q)$ – the best-shot – out of a total of $f_i(a_i)$ draws. In an innovation contest, H_i can be interpreted as the distribution of the quality of a single idea and f_i as the number of ideas. Formally, however, there is no reason to restrict f_i to take integer values but it must be non-negative, $f_i(a_i) \geq 0$.

The setting is inspired by Fullerton and McAfee (1999). However, they and the ensuing literature assume that $H_i(q) = H(q)$ for all $i \in N$. In words, all agents have ideas that are equally good ex ante but some agents may have more ideas than others. It is straightforward to show that using (4) in (3) produces (1). This is intuitive. After all, agent i makes $f_i(a_i)$ draws from $H(\cdot)$ out of a total of $\sum_{i\in N} f_j(a_j)$ draws. Each draw has an equal chance of being the highest draw, thus yielding the CSF in (1). Note that if $f_i(0) = 0$ then $G_i(q_i|0)$ is degenerate with all mass at $q_i = \underline{q}_i = \underline{q}$. Thus, using the specified tie-breaking rule, agents win with equal probability if all actions are zero. Since $f_i(a_i)$ can be thought of as the number of ideas, it is often sensible to assume that $f_i(a_i)$ is increasing and concave, or $f_i'(a_i) > 0 \geq f_i''(a_i)$. Then, $p_i(a_1, ..., a_n)$ is concave in a_i.

Things are more complicated if the H_i's are allowed to be heterogenous. For instance, if H_i first-order stochastically dominates H_j for all $j \neq i$ then (1) is only a lower bound on agent i's probability of winning. The reason is that agent i's ideas are better ex ante than his rivals' ideas. Such a setting thus cannot be analyzed using (1), but it turns out to be amenable to the approach suggested in this paper. Henceforth, the "best-shot model" refers to (4) with potentially heterogenous H_i's. The special case in which all H_i's are identical is referred to as the Fullerton and McAfee (1999) model. $\square$

Returning to the general model, note that (3) requires only that it can be identified whose performance is the highest. In this paper, however, it is assumed that all performances are observed. A biased contest is then one in which the winner is not necessarily the agent with the highest performance.

3 Contests as moral hazard problems

A contest elicits effort from agents. Hence, designing a contest is at heart a moral hazard problem. This preliminary section begins by first describing the manner in which the contest can be manipulated and compares it to other approaches in the literature. Some important technical assumptions are also introduced and discussed. Finally, the designer's objective function is discussed.

3.1 Assignment rules

One way of viewing (2) is as a family of black-box CSFs that the designer can choose from. The choice to focus on this family is arbitrary but it is unclear what the "most reasonable" family is. This is a fundamental problem with any black-box approach: Biasing the contest is akin to "shaking" the black box, but it is not obvious how to model the consequences of this intervention.

On the other hand, the underlying stochastic-performance structure means that one can now think of the CSF not as a black box but rather, borrowing a term from computer engineering, as a "white box" because the internal workings of the system are known. The hard-wired components are described by $G_i(q_i|a_i)$, on top of which is a program that identifies the winner as being the top performer.

The white-box approach provides clear direction for how to approach contest design. Optimal design boils down to asking how the box can best be "hacked", i.e. how to reprogram the mapping from performances to outcomes.

Thus, the designer constructs n functions, $P_i(\mathbf{q})$, $i \in N$, that describe the winning probability of each agent, contingent on performances. Here, $\mathbf{q} = (q_1, q_2, ..., q_n)$ denotes the profile of performances. The only constraints are that $P_i(\mathbf{q}) \geq 0$, $i \in N$, and $\sum P_i(\mathbf{q}) \leq 1$. Together, $P_i(\mathbf{q})$, $i \in N$, define an "assignment rule." It is often convenient to write $P_i(\mathbf{q})$ as $P_i(q_i, \mathbf{q}_{-i})$, where $\mathbf{q}_{-i}$ denotes the vector of performances of agent i's rivals. It is explicitly assumed that the designer can credibly and fully commit to any feasible assignment rule. In comparison, it is not always clear what is implicitly or explicitly assumed in this regard in work that relies on (2). How realistic the assumption is depends on the application. Che and Gale (2003) consider a research contest in which the quality of the innovation is not verifiable. Then, the designer can obviously not commit to any arbitrary assignment rule.

Let $\mathbf{a}_{-i}$ denote the vector of actions by agent i's rivals. Given $\mathbf{a}_{-i}$, agent i's expected utility from action a_i is

$$U_i(a_i, \mathbf{a}_{-i}) = v_i \int \left(\int P_i(q_i, \mathbf{q}_{-i}) g_i(q_i|a_i) dq_i \right) \prod_{j \neq i} g_j(q_j|a_j) d\mathbf{q}_{-i} - a_i, \tag{5}$$

since performances are statistically independent. For a profile of actions $\mathbf{a} = (a_1, a_2, ..., a_n)$ to be implementable, it must of course constitute a Nash Equilibrium of the contest game. The point is that this equilibrium can be manipulated by making changes to the assignment rule. Attention is restricted to pure strategy implementation throughout.

3.2 Related contracting problems

Contest design shares some similarities with contracting in team moral hazard problems, à la Holmström (1982). In the latter, the principal designs wage schedules, $w_i(\mathbf{q})$, with complete freedom. In contest design as defined here, there are no monetary transfers (at least not beyond the prize, which may or may not be monetary). Nevertheless, agent i cares about $v_i P_i(\mathbf{q})$ in a contest in much the same way as he cares about $w_i(\mathbf{q})$ in Holmström's (1982) setting. The constraint that $P_i(\mathbf{q}) \geq 0$ is akin to a limited liability constraint. The constraint that $\sum P_i(\mathbf{q}) \leq 1$ is similar to a budget constraint. Thus, there is less design flexibility in contests than in classic team moral hazard problems. This has important consequences. In Holmström (1982), $w_i(\mathbf{q})$ depends only on q_i if all signals are independent, as is the case in the present model. Such wage schedules may of course not be feasible if there is a budget constraint. It is then necessary to make pay contingent on the performance of other agents as well to make sure the budget is not broken. For similar reasons of feasibility it is generally optimal to let $P_i(\mathbf{q})$ depend on the entire profile $\mathbf{q}$ in the contest setting.

Rank-order tournaments as in Lazear and Rosen (1981) are essentially team moral hazard problems where wage schedules are restricted to take very particular functional forms. In particular, wages can take one of n values and they must be allocated in order of agents' performances. Thus, rank-order tournaments can also be viewed as special kinds of contests with n prizes in which the assignment rule is restricted but where the values of the prizes are design instruments. Lazear and Rosen (1981) consider an extension to two-agent tournaments in which handicapping is used to determine the winner. This takes the specific form of an additive bias applied to performances and so this still places restrictions on $P_i(\mathbf{q})$. See also Fain (2009).

3.3 The contest environment

Unless explicitly mentioned, no functional form is imposed on the distributions $G_i(q_i|a_i)$. Thus, the aim is to analyze contests in some generality. However, there are technical road blocks. In this first paper, the idea is to make use of the simplest possible techniques in order to focus on conceptual and economic insights. Thus, the analysis relies on the standard first-order approach known from classic moral hazard problems. The validity of this approach places technical restrictions on $G_i(q_i|a_i)$ and thereby limits the set of $G_i(q_i|a_i)$ for which the optimal design is characterized.[6] It is, however, important to note that this is a purely technical

[6] For instance, the assumption that the support of q_i is independent of the action for $a_i > 0$ simplifies the incentive compatibility constraint and is standard in the moral hazard literature.

problem. This is in contrast to the deeper conceptual limitations that are implicit in restricting attention to $G_i(q_i|a_i)$ that yield (1).

It will be assumed from now on that actions are continuous and that $g_i(q_i|a_i)$ is differentiable with respect to a_i when $a_i > 0$. Agent i's likelihood-ratio,

$$L_i(q_i|a_i) = \frac{1}{g_i(q_i|a_i)} \frac{\partial g_i(q_i|a_i)}{\partial a_i},$$

plays an important role. A common assumption is that $L_i(q_i|a_i)$ is weakly increasing in q_i. For expositional simplicity, this paper assumes that $L_i(q_i|a_i)$ is strictly increasing in q_i. This will be referred to as the *monotone likelihood-ratio property* (MLRP). The MLRP implies that higher actions make lower performances less likely. Thus, $G_i(q_i|a_i)$ is strictly decreasing in a_i whenever q_i is interior. Similarly, agent i's expected performance, $\mathbb{E}[q_i|a_i]$, is strictly increasing in a_i.

In the standard contracting literature, the role of the MLRP is to ensure that wage schedules are monotonic in signals. It plays a similar role here. Agent i's assignment $P_i(q_i, \mathbf{q}_{-i})$ is said to be monotonic if it is non-decreasing in q_i. While agent i's expected utility clearly depends on the properties of $P_i(q_i, \mathbf{q}_{-i})$, a standard technique can be applied whenever $P_i(q_i, \mathbf{q}_{-i})$ is monotonic in equilibrium. Specifically, Rogerson (1985) combines the MLRP with a *convexity of the distribution function condition* (CDFC) that assumes that $G_i(q_i|a_i)$ is convex in a_i for all q_i. The CDFC implies that the term in the parenthesis in (5) is concave in a_i for any monotonic $P_i(q_i, \mathbf{q}_{-i})$; the easiest way to see this is by using integration by parts. Thus, agent i's expected utility is concave in a_i, given $\mathbf{a}_{-i}$. Consequently, the first-order condition identifies a best response.

The best-shot model in (4) has the MLRP whenever $f_i'(a_i) > 0$ and it likewise satisfies the CDFC whenever $f_i''(a_i) \leq 0$. Indeed, Rogerson's (1985) leading example is precisely a special case of (4). Thus, the current paper concentrates on environments where the MLRP and the CDFC hold. It can be verified that the approach extends to Hirschleifer and Riley's (1992) two-agent model when $f_i(a_i) = a_i$, despite the fact that such a model violates the CDFC.

3.4 Objective and welfare functions

Returning to contest design, the optimal assignment rule generally depends on the designer's objective function. There are at least two conceptually very distinct ways to think about objective functions. First, the designer may care directly about actions. Perhaps in part for historical reasons, this is the prevalent assumption in contest theory, either implicitly or explicitly. Recall that Tullock (1980) examined rent-seeking contests. Here, effort is pure waste and it thus makes sense from a welfare perspective to understand the size of the loss, measured by $\sum_{i\in N} a_i$. However, many applications of (1) since Tullock consider settings where higher actions benefit society or possibly some contest designer. Hence, it is common to measure welfare by $\sum_{i\in N} a_i$ and, when the contest design is endogenous, to seek to maximize this. More generally, the welfare or objective function can be captured by a benefit function $B(\mathbf{a})$ that depends only on the action profile.

Second, the designer may care directly only about performances. Then, actions are important only via their impact on performances. For instance, consider a promotion contests among salesmen. Here, the employer is presumably not directly interested in the salesmens' efforts, $\sum_{i\in N} a_i$, but rather in the expected total volume of sales, $\mathbb{E}[\sum_{i\in N} q_i|\mathbf{a}]$. More generally, the designer has some Bernoulli utility function $\pi(\mathbf{q})$ that depends only on performances. Expected utility is then $B(\mathbf{a}) = \mathbb{E}[\pi(\mathbf{q})|\mathbf{a}]$, which is again a function of the action profile $\mathbf{a}$.

In either case, the expected utility function is written as $B(\mathbf{a})$. The motivation is different however, and it is important conceptually to point out that $\sum_{i\in N} a_i$ and $\mathbb{E}[\sum_{i\in N} q_i|\mathbf{a}]$ may generally be quite different. For instance, in the Fullerton and McAfee (1999) model, $\mathbb{E}[\sum_{i\in N} q_i|\mathbf{a}]$ depends on $H(q)$. If $H(q) = q^\alpha$, $\alpha > 0$, $q \in [0, 1]$, then

$$\mathbb{E}[\textstyle\sum_{i\in N} q_i|\mathbf{a}] = \textstyle\sum_{i\in N} \dfrac{\alpha f_i(a_i)}{1 + \alpha f_i(a_i)}. \tag{6}$$

If the designer cares about performances rather than actions per se, then (6) is evidently of much more interest than $\sum_{i\in N} a_i$. Note that the two are typically not maximized at the same action profile.

However, assuming that the designer's expected utility depends only on $\mathbf{a}$ is not without loss of generality. The reason is that such utility functions are *assignment independent*: The designer's utility directly or indirectly depends only on actions but it is independent of the assignment rule itself. This is not the case in an innovation contest if the designer is restricted to implementing the winning project, even if one of the losing contestants has a better project. Then, the expected quality of the winning project is

$$\int \left(\sum_{i \in N} P_i(\mathbf{q}) q_i \right) \prod_{i \in N} g_i(q_i|a_i) d\mathbf{q},$$

which evidently depends on the assignment rule directly.

Thus, there are a large number of potentially interesting objective function. Which objective function applies depends on the application. This paper focuses on objective functions that are assignment independent. In many applications, it is reasonable to assume that $B(\mathbf{a})$ is strictly increasing in all arguments. An objective function that is assignment independent and strictly increasing will be said to satisfy *Assignment Independence and Monotonicity*, abbreviated AIM. Not only is this a reasonably large class and a good place to start, it also turns out that optimal contest design is qualitatively similar for all such objective functions.

If $B(\mathbf{a})$ takes the form $B(\mathbf{a}) = \mathbb{E}[\pi(\mathbf{q})|\mathbf{a}]$, then the MLRP implies AIM if π is strictly increasing in each argument. The assumption that π is *strictly* increasing is sufficient but not necessary. Assume that all agents' performances share the same support. Then, $\pi(\mathbf{q}) = \max\{q_1, q_2, .., q_n\}$ also gives an objective function that satisfies AIM. This Bernoulli utility function applies when the designer only cares about the best performance, even though this may or may not equal the winner's performance. For instance, a firm may pursue the best product design proposed by a disparate group of in-house developers, yet may at the same time chose to promote a developer whose own design was inferior to handle the product launch. Likewise, if no distribution is degenerate even at the zero action, then $\pi(\mathbf{q}) = \min\{q_1, q_2, .., q_n\}$ also implies AIM.[7] This objective function may apply in a team product design problem where the quality of the worst component determines the overall value or longevity of the product.

[7] The case where distributions are degenerate at the zero action is not too different. Here, it can never be optimal to induce $a_i = 0$ for any $i \in N$ given that $\pi(\mathbf{q}) = \min\{q_1, q_2, .., q_n\}$. Then, AIM is satisfied on the subset of relevant actions, i.e. those for which $a_i > 0$ for all $\in i \in N$.

4 Optimal contest design

Grossman and Hart (1983) propose a two-step procedure to the classic moral hazard problem. The first step derives the cost-minimizing way of implementing any given action. The second step then solves for the optimal action to implement, taking into account both implementation costs and the benefits to the principal.

The procedure needs modification in the current setting. Assume that the objective function satisfies AIM. Then, given the prize is fixed and exogenous, implementation costs are invariant to the action profile and independent of the assignment rule.[8,9] The problem is instead that the restriction to using a contest implies that not all action profiles can be implemented. Thus, the natural first step is to ask which action profiles can be implemented. The second step then implements the action profile that is most beneficial to the designer. If the objective function satisfies AIM, the optimal action must be on the frontier of the implementable set of actions.

This section breaks the analysis of assignment rules into a few parts. First, the "feasible set" of actions that can be elicited from any individual agent is characterized. For any implementable action, there are generally a multitude of assignment rules that are incentive compatible. However, the $P_i(\mathbf{q})$ that induces the highest possible action from agent i is unique.

The second part considers contests in which the prize must be allocated and concentrates on action profiles where $a_i > 0$ for all i. Again, incentive compatible assignment rules are not unique for most implementable action profiles. However, for any action profile that is along the frontier of the feasible set, there is in fact a unique incentive compatible contest design. The structure of the assignment rule is similar for all profiles along the frontier. The implication is that as long as the objective function satisfies AIM, the fundamental structure of the optimally designed contest is uniquely characterized.

The third part considers contests in which some agents may be inactive and contests in which it is possible to ration the prize. Once again, all objective functions satisfying AIM yield assignment rules with the same fundamental structure.

[8] Even though the prize is fixed it may be valued differently by different agents, or $v_i \neq v_j$.

[9] In some applications it may be costly to allocate the prize. The designer may thus decide to not assign the prize. The objective function does not satisfy AIM in such cases.

4.1 Maximal individual effort

Given (5), the marginal return to a small increase in a_i is

$$\frac{\partial U_i(a_i, \mathbf{a}_{-i})}{\partial a_i} = v_i \int \left(\int P_i(q_i, \mathbf{q}_{-i}) L_i(q_i|a_i) g_i(q_i|a_i) dq_i \right) \prod_{j \neq i} g_j(q_j|a_j) d\mathbf{q}_{-i} - 1. \quad (7)$$

Since the expected value of $L_i(q_i|a_i)$ is zero, it follows from the MLRP that $L_i(q_i|a_i)$ is strictly negative for small q_i and strictly positive for large q_i. It is clear that (7) is maximized if the prize is assigned to agent i if and only if $L_i(q_i|a_i)$ is positive. When $L_i(q_i|a_i)$ is positive, a marginal increase in a_i makes it more likely that a performance close to q_i is realized. There is no better carrot than promising the agent the prize for such performances and no better stick than to deny him the prize for performances that become less likely if his action increases.

Let $\widehat{q}_i(a_i)$ denote the unique value of q_i for which $L_i(q_i|a_i) = 0$. Now fix some target action, a_i^t, that the designer may wish to implement. When evaluated at $a_i = a_i^t$, (7) is thus maximized with an assignment rule that has the property that

$$P_i(q_i, \mathbf{q}_{-i}) = \begin{cases} 1 & \text{if } q_i \geq \widehat{q}_i(a_i^t) \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

Any assignment rule that takes a form such as that in (8) will be said to be a *threshold rule for agent i.* In a promotion contest, for example, agent i might be the "heir apparent" who is destined to win the promotion unless his performance is a conspicuous failure. This puts maximal pressure on the agent to ensure that his performance lives up to expectations.

Note that the threshold rule is independent of $\mathbf{q}_{-i}$. Hence, agent i's incentives are the same regardless of the actions taken by other agents. Likewise, it leaves unspecified to whom the prize is assigned if agent i fails to meet the threshold. This is of course irrelevant from agent i's perspective, but it does potentially impact the equilibrium actions of other agents.

The threshold rule in (8) is useful because by construction it maximizes the first derivative in (7) when evaluated at $a_i = a_i^t$. Hence, if the threshold rule leads (7) to take a negative value, then there is no assignment rule that can feasible satisfy the first-order condition. Then, the target action a_i^t simply cannot be implemented. Hence, it is necessary for implementability that (7) is non-negative at a_i^t when the threshold rule is used.

Given (8), agent i's expected utility from some action a_i is

$$v_i(1 - G_i(\widehat{q}_i(a_i^t)|a_i)) - a_i. \tag{9}$$

Hence, following the above argument, implementability necessitates that the derivative with respect to a_i is non-negative when evaluated at $a_i = a_i^t$, or

$$-\frac{\partial G_i(\widehat{q}_i(a_i^t)|a_i)}{\partial a_i}_{|a_i=a_i^t} \geq \frac{1}{v_i}. \tag{10}$$

As noted earlier, the MLRP implies that the left-hand side is strictly positive.

Lemma 1 *Given the MLRP and some interior target a_i^t, there exists an assignment rule that satisfies agent i's first-order condition at $a_i = a_i^t$ if and only if (10) is satisfied. Moreover, the $P_i(q_i, \mathbf{q}_{-i})$ function that satisfies agent i's first-order condition is (essentially) unique if and only if (10) is binding.*[10]

Proof. See the Appendix. ■

Let $\overline{a}_i$ denote the highest possible value of a_i^t for which (10) holds, assuming such a value exists. By Lemma 1, $\overline{a}_i$ is a candidate for the highest possible action that agent i can be induced to take. From now on, it will be assumed that $\overline{a}_i$ exists and is strictly positive. This is without loss of generality as otherwise the set of agents can be redefined to contain only those agents that can be induced to take a positive action.

Lemma 1 relies on the first-order condition. However, since threshold rules are monotonic, the first-order condition is sufficient when the CDFC holds.

Proposition 1 *Assume G_i satisfies the MLRP and CDFC. Then, any interior a_i^t for which (10) holds can be implemented. In particular, $\overline{a}_i$ is implementable.*

Proof. The proof of Lemma 1 describes a threshold rule that satisfies the first-order condition at $a_i = a_i^t$. By concavity, the agent has no incentive to deviate from this action. ■

[10] Here, $P_i(q_i, \mathbf{q}_{-i})$ is "essentially unique" because changes on a set of performances of measure zero are irrelevant.

4.2 Contests with no inactive agents and no rationing

The remainder of the paper considers all agents together, as is required for optimal contest design. In general, it may be optimal to threaten to not allocate the prize. However, such rationing is rarely considered or allowed in the current literature on contest design. Thus, attention is for now restricted to the optimal design of contests subject to the restriction that the prize must be allocated. Formally,

$$\textstyle\sum_{i \in N} P_i(\mathbf{q}) = 1.$$

The assumption is realistic in many applications. For instance, a CEO must be found eventually.

Let $\mathcal{A}_N$ denote the set of action profiles $\mathbf{a}$ that can be implemented subject to the no-rationing constraint. Given the MLRP and the CDFC, this contains action profiles where $a_i = \overline{a}_i$ for some $i \in N$ and $a_j = 0$ for all $j \neq i$. To implement this, apply the threshold rule to agent i and if agent i is unsuccessful then give the prize at random to one of the other agents.

This subsection describes the portion of the frontier where *all* agents are *active*, or $a_i > 0$ for all i. It turns out that allowing for inactive agents, or $a_i = 0$ for some i, is similar to allowing for rationing. These possibilities are taken up in the next subsection.

Fix some agent j and decompose an interior action profile $\mathbf{a}$ into $(a_j, \mathbf{a}_{-j})$. If $\mathbf{a}$ is on the frontier of $\mathcal{A}_N$, then it must necessarily hold that no higher value of a_j can be implemented given $\mathbf{a}_{-j}$. Thus, the question is: Given a fixed $\mathbf{a}_{-j}$, what is the highest implementable value of a_j? To answer this, recall that the first-order conditions are necessary for all agents other than agent j since $\mathbf{a}_{-j}$ is interior. Without loss of generality let $j = 1$. Extending the approach that gave Lemma 1 (where $\mathbf{a}_{-1}$ was essentially ignored), fix $(a_1, \mathbf{a}_{-1})$ and consider the problem

$$\begin{aligned} &\max_{\{P_i(q) \in [0,1]\}_{i=1}^n} \frac{\partial U_1(a_1, \mathbf{a}_{-1})}{\partial a_1} \\ st \quad &\frac{\partial U_i(a_i, \mathbf{a}_{-i})}{\partial a_i} = 0 \quad i \in N \backslash \{1\} \\ &\textstyle\sum_{i \in N} P_i(\mathbf{q}) = 1 \text{ for all } \mathbf{q}. \end{aligned} \tag{11}$$

Let $V(a_1, \mathbf{a}_{-1})$ denote the maximum-value function of this problem. Note that $V(a_1, \mathbf{a}_{-1})$ plays the same role as (10) does in Lemma 1 and Proposition 1. If $V(a_1, \mathbf{a}_{-1}) < 0$ then it is impossible to construct a contest that satisfies agent 1's first-order condition and the action profile $(a_1, \mathbf{a}_{-1})$ cannot be implemented.[11] Given the MLRP and the CDFC are satisfied as in Proposition 1, it will be shown that there is an essentially unique assignment rule that implements the action profile if $V(a_1, \mathbf{a}_{-1}) = 0$.

Let μ_i denote the multiplier to the incentive constraint for agent i, $i \in N\backslash\{1\}$. Define $\mu_1 = 1$. It is convenient to write the last constraint as

$$\left(\sum_{i\in N} P_i(\mathbf{q}) - 1\right) \prod_{i\in N} g_i(q_i|a_i) = 0.$$

Let $\eta(\mathbf{q})$ denote the multiplier to this constraint when the performance profile is $\mathbf{q}$. The Lagrangian can then be written as

$$\begin{aligned} \mathcal{L} &= \sum_{i\in N} \mu_i \frac{\partial U_i(a_i, \mathbf{a}_{-i})}{\partial a_i} + \int \eta(\mathbf{q}) \left(\sum_{i\in N} P_i(\mathbf{q}) - 1\right) \prod_{i\in N} g_i(q_i|a_i) d\mathbf{q}. \\ &= \int \left(\sum_{i\in N} P_i(q_i, \mathbf{q}_{-i}) \left[\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q})\right] - \eta(\mathbf{q})\right) \prod_{i\in N} g_i(q_i|a_i) d\mathbf{q} - \sum_{i\in N} \mu_i \end{aligned}$$

Maximizing $\mathcal{L}$ pointwise for any given $\mathbf{q}$ by appropriately choosing the assignment rule is equivalent to maximizing

$$\sum_{i\in N} P_i(q_i, \mathbf{q}_{-i}) \left[\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q})\right] \tag{12}$$

subject to the feasibility constraints. This is done by letting $P_i(\mathbf{q}) = 1$ if

$$\mu_i v_i L_i(q_i|a_i) > \max_{j\neq i}\{\mu_j v_j L_j(q_j|a_j)\} \tag{13}$$

and $P_i(\mathbf{q}) = 0$ if the inequality is reversed. It is as if agent i earns a score of $\mu_i v_i L_i(q_i|a_i)$ and the agent with the highest score wins. Ties occur with probability zero and it is irrelevant how they are broken. This rule is fairly intuitive. As discussed in the previous subsection, the power of the incentives facing agent

[11] Similarly, if the problem has no solution then it is impossible to simultaneously satisfy the first-order conditions of agents $2, ..., n$. Thus, the action profile cannot be implemented.

i are determined by the size of $v_i L_i(q_i|a_i)$ when he is assigned the prize. Hence, (13) describes how to optimally balance incentives across agents. Note that as $\mu_j \to 0$ for all $j \neq i$, the assignment rule converges to a threshold rule for agent i, such that he wins if and only if his likelihood-ratio is positive.

Given the MLRP, a standard argument shows that $\mu_i > 0$. Then, agent i is more likely to win the higher q_i is. In other words, the assignment rule is monotonic. The CDFC then implies that agent i's utility is concave in his action.

Proposition 2 *Assume that G_i satisfies the MLRP and the CDFC for all $i \in N$, and assume that the prize must be allocated. Fix an action profile $\mathbf{a}$ in which all agents are active. If $\mathbf{a}$ is on the frontier of $\mathcal{A}_N$, then there is an essentially unique assignment rule that implements it. This assignment rule is described by (13), where $\mu_i \in (0, \infty)$ are endogenously determined constants, $i \in N$.*

Proof. See the Appendix. ■

If the objective function satisfies AIM then the optimal action to implement must be on the frontier of $\mathcal{A}_N$. The frontier, however, consists of action profiles where all agents are active and other action profiles where only a subset are active. As long as the solution entails only active agents then the optimal contest design must be given by the rule described in (13). Note that this structure is independent of the exact functional form of the objective function $B(\mathbf{a})$. Action profiles with inactive agents are considered in the next section. This leads to a minor and fairly straightforward modification of (13).

Conversely, imagine that the contest designer implements an action profile in which all agents are active but with an assignment rule that does not take the form in (13). Then, the action profile is not on the frontier of $\mathcal{A}_N$, since such action profiles can only be implemented using (13). In this case, then, the contest designer does not have an objective function that satisfies AIM.

4.3 Inactive agents and rationing

There are many different assignment rules that induce an agent to be inactive. For instance, any rule in which $P_i(q_i, \mathbf{q}_{-i})$ is independent of q_i induces $a_i = 0$. Inducing someone to be inactive is easy. Hence, there is generally no unique assignment rule that implements an action profile with inactive agents.

However, an alternative way of thinking of the problem is to just restrict attention to the active agents and ignore the existence of the inactive ones. In other words, delete from the performance profile $\mathbf{q}$ all the performances of the inactive agents. The assignment rule may then be unique as a function of the performances of the active agents.[12] If none of the active agents are assigned the prize, then it is simply given away to one of the inactive agents at random. The inactive agents can not influence the assignment and so they indeed have no incentive to become active.

Let $\mathcal{N}$ denote the subset of active agents in the population N of agents. Winning probabilities are not restricted to sum to one among the agents in $\mathcal{N}$. With some abuse of notation, $\mathbf{a}$ and $\mathbf{q}$ now denote the action profile and the performance profile, respectively, of the active agents. By definition $a_i > 0$ for all $i \in \mathcal{N}$ since $\mathcal{N}$ describes the active agents. Likewise, $a_i = 0$ if $i \notin \mathcal{N}$.

The problem of selecting the optimal set of active contestants is in the background. It is ignored for now but taken up again later. The primary focus for now remains more narrowly on identifying a robust structure of optimal contest.

Adding the possibility of rationing is similar to adding one or more extra agents that are inactive in equilibrium. If the prize is withheld from the active agents, it can just be dumped with an inactive agent. Thus, even if rationing may appear hard to commit to, it is possible in a contest with $n+1$ agents to reproduce the optimal action profile in an n agent contest that allows for rationing. From the point of view of the original n agents, the threat of giving the prize to the additional agent is equivalent to threaten to withhold the prize entirely. Hence, there are two ways to interpret a contest with a set $\mathcal{N}$ of active agents where winning probabilities are not restricted to sum to one: (1) The prize must be allocated but there are inactive agents that occasionally receive the prize, or (2) the designer has the power to ration or withhold the prize. Let $\mathcal{A}_{\mathcal{N}}^{R}$ denote the set of action profiles that can be induced among the set $\mathcal{N}$ of active agents when rationing is allowed. Again, the frontier of this set is described.

It is straightforward to modify Proposition 2 to allow for rationing. In (11) –

[12] This approach is without loss of generality. Given q_i, agent i only cares about the expectation of $P_i(q_i, \mathbf{q}_{-i})$ over $\mathbf{q}_{-i}$. The set of performances of the inactive agents can then always be replaced by some randomization device.

assuming for the sake of argument that agent 1 is active – the last equality constraint must be replaced by an inequality constraint since winning probabilities need not sum to one, and $\mathcal{N}$ replaces N. This does not materially change the expression in (12). Hence, if the prize is assigned, it is once again assigned to the agent with the highest $\mu_i v_i L_i(q_i|a_i)$.

Proposition 3 *Assume that G_i satisfies the MLRP and the CDFC for all $i \in \mathcal{N}$. Assume that the prize need not be assigned with probability one among $\mathcal{N}$. If a given action profile $\mathbf{a}$ is on the frontier of $\mathcal{A}_{\mathcal{N}}^R$, then there is an essentially unique assignment rule that implements it. Here, $P_i(\mathbf{q}) = 1$ if*

$$\mu_i v_i L_i(q_i|a_i) > \max\{0, \max_{j \neq \mathcal{N} \setminus \{i\}} \{\mu_j v_j L_j(q_j|a_j)\}\} \tag{14}$$

and $P_i(\mathbf{q}) = 0$ if the inequality is reversed, where $\mu_i \in (0, \infty)$ are endogenously determined constants, $i \in \mathcal{N}$.

Proof. See the Appendix ■

Proposition 3 combines Propositions 1 and 2. Using the logic that lead to the threshold rule used in Lemma 1 and Proposition 1, assigning the prize when the likelihood-ratio is negative weakens incentives. Thus, the prize must be withheld in such cases. Moreover, to balance incentives across agents, Proposition 2 reveals how scores should be compared. The rule in (14) have all these features.

To illustrate, consider the special case where $N = \{1, 2\}$ and $\mathcal{N} = \{1\}$. Then, the frontier of $\mathcal{A}_{\mathcal{N}}^R$ is simply $\overline{a}_1$; (14) is just the optimal threshold rule from Lemma 1 and Proposition 1. This can be thought of as describing the action profile $(a_1, a_2) = (\overline{a}_1, 0)$, where agent 2 is used as the threat to enforce rationing.

More generally, assume that $\mathcal{N}$ is a proper subset of N, i.e. that there is one or more inactive agents. Compare the frontier of $\mathcal{A}_{\mathcal{N}}^R$ with the part of the frontier of $\mathcal{A}_N$ where all agents are active. The latter describes an action profile where $a_i > 0$ for all $i \in N$. The former specifies an action profile of the set $\mathcal{N}$ of active agents, but it is understood that $a_i = 0$ for all $i \in N \backslash \mathcal{N}$ at the same time. Thus, the frontier of $\mathcal{A}_{\mathcal{N}}^R$ is the part of the frontier of $\mathcal{A}_N$ where a subset of agents are inactive. In this way Proposition 2 and 3 together describe the entire frontier of $\mathcal{A}$. Once again, the structure of the optimal contest is remarkably robust. That is, the assignment rule is determined by a comparison of scaled likelihood-ratios for all objective functions that satisfy AIM.

4.4 Selecting the active agents

Determining the optimal set of active agents, $\mathcal{N}$, is generally complicated. First, agents may have different distribution functions. To rule this out, assume now that $G_i(q_i|a_i) = G(q_i|a_i)$ for all $i \in N$. Second, the designer may care more about some agents compared to other agents. To eliminate this complication, assume that the benefit function is anonymous. Under these assumptions, the active agents are those with the highest valuations. That is, if agent i is induced to be active and agent j is induced to be inactive, it must be the case that $v_i \geq v_j$.

The reason is simple. Suppose to the contrary that $v_i < v_j$ but that agent i is active and agent j is not. Let agent i play the role of agent 1 in the problem (11). If a_i is on the boundary of $\mathcal{A}^R_{\mathcal{N}}$, then the objective function in (11) is zero. However, replacing agent i with agent j leads to a larger objective functions. This implies that agent j can be induced to take a higher action than a_i, even while keeping all other agents at the same level of effort.

This observation is intuitive. Simply put, it is easier to induce high actions from agents who value the prize very highly. In fact, Fu and Wu (2020) obtain the same result in their model based on (2). However, it should be contrasted to the influential "exclusion principle" due to Baye, Kovenock, and de Vries (1993). In their model the CSF is exogenously given as an all-pay auction CSF. In that setting, it may be optimal to exclude agents with high valuations. In the all-pay auction, such agents deter other agents from exerting much, or any, effort. In the current model, in contrast, the CSF is endogenous.

5 The best-shot model revisited

This section utilizes the best-shot model to illustrate some of the main results. This also makes it possible to return to the issue of whether (2) and the conclusions stemming from it are compatible with the stochastic performance model.

5.1 Optimal assignment rules

From (13), the likelihood-ratios play a crucial role in optimal contest design. Generally speaking, likelihood-ratios are non-linear functions and the optimal assignment rule is therefore typically a complicated function of performance levels.

In the best-shot model, for any action $a_i > 0$,

$$L_i(q_i|a_i) = \frac{f_i'(a_i)}{f_i(a_i)} + f_i'(a_i) \ln H_i(q_i).$$

Note that the best-shot model has the special property that $L_i(q_i|a_i) \to -\infty$ as $q \to \underline{q}_i$. Therefore, (13) and (14) imply that any agent whose performance is $\underline{q}_i$ wins with probability zero when an action profile along the frontier of $\mathcal{A}_N$ or $\mathcal{A}_N^R$ is implemented.

To better understand the optimal assignment rule along the frontier of $\mathcal{A}_N$ or $\mathcal{A}_N^R$, it is necessary to have a closer look at the adjusted scores. If the equilibrium action is a_i^* then agent i's adjusted score is

$$\begin{aligned} s_i(q_i) &= \mu_i v_i L_i(q_i|a_i^*) \\ &= \tau_i \left(1 + \ln G_i(q_i|a_i^*)\right), \end{aligned} \tag{15}$$

where

$$\tau_i = \mu_i v_i \frac{f_i'(a_i^*)}{f_i(a_i^*)} > 0.$$

Hence, the optimal assignment rule can be implemented as follows. First, q_i is translated into the quantile where it sits in the equilibrium distribution $G_i(q_i|a_i^*)$. This intermediate score then undergoes a monotonic transformation before being multiplied by an endogenous and identity-dependent constant, τ_i. The winner is the agent with the highest final score.

The score is a non-linear function of performance. This is a general property that extends beyond the best-shot model. Thus, it is generally not sufficient to compare affine transformations of performances.

Returning to the best-shot model, note that if agents i and j perform equally well given what is expected of them – i.e. they perform at the same quantiles, or $G_i(q_i|a_i^*) = G_j(q_j|a_j^*)$ – then agent i beats agent j if $\tau_i > \tau_j$. Hence, in equilibrium, the ex ante winning probabilities are ordered the same way as the τ_i's. In this sense, τ_i is a measure of how favorable the contest is to agent i.

Given a vector $\boldsymbol{\tau}$ that lists all τ_i's for the active agents, $i \in \mathcal{N}$, it is in principle possible to derive the probability that agent i wins for any given action profile $\mathbf{a}$. In other words, the endogenous CSF can be characterized. To illustrate, the CSF for the least favored agent is derived next, under the assumptions that all agents are active and that the prize must be allocated. These assumptions are made to make the comparison with (2) easier.

Proposition 4 *Consider the best-shot model with $f_j'(a_j) > 0 \geq f_j''(a_j)$ for all $j \in N$ and fix an action profile $\mathbf{a}^*$ on the frontier of $\mathcal{A}_N$ in which all agents are active, or $a_j^* > 0$ for all $j \in N$. If $\tau_i \leq \tau_j$ for all $j \in N$, then agent i wins with probability*

$$\widehat{p}_i(\mathbf{a}|\boldsymbol{\tau}) = \left(\prod\nolimits_{j \in N} e^{\frac{(\tau_i - \tau_j) f_j(a_j)}{\tau_j f_j(a_j^*)}} \right) \frac{f_i(a_i)}{\sum_{j \in N} \frac{\tau_i f_i(a_i^*)}{\tau_j f_j(a_j^*)} f_j(a_j)}, \tag{16}$$

for any action profile where $a_i > 0$.

Proof. See the Appendix. ■

The last term in (16) is similar to (2) when only handicaps are used. However, since $\tau_i \leq \tau_j$ for all $j \in N$, the first term is less than one and agent i thus wins less often than what (2) would suggest. As a consistency check, note that if $\tau_i = \tau_j$ for all $j \in N$ then $\widehat{p}_i(\mathbf{a}^*|\boldsymbol{\tau}) = \frac{1}{n}$ and all agents win with equal probability in equilibrium.

Importantly, the first term in (16) depends on the action profile. Thus, (16) is *not* a ratio-form or generalized-lottery CSF (except in the special case where $\tau_i = \tau_j$ for all $j \in N$). Applying Proposition 4 to the Fullerton and McAfee (1999) model implies that even if the unbiased contest is a generalized lottery contest, this need no longer hold when the contest design is endogenous.

5.2 Optimal action profiles

The designer can always transform q_i into the quantile $\widetilde{q}_i = H_i(q_i)$ and use this as the basis for contest design. Given agent i's action, the distribution of $\widetilde{q}_i$

is $\widetilde{q}_i^{f_i(a_i)}$, $\widetilde{q}_i \in [0,1]$, independently of what $H_i(q_i)$ is. It follows that the set of implementable action profiles is independent of the distributions of ideas. However, it is important to note that the optimal action profile to implement may be sensitive to $H_i(q_i)$ when the designer has an objective function of the form $B(\mathbf{a}) = \mathbb{E}[\pi(\mathbf{q})|\mathbf{a}]$. The reason is that the expected value generally speaking depends on the distribution of performances.

A natural place to begin is by utilizing the structure in the best-shot model to more succinctly characterize the highest possible implementable action of agent i, $\overline{a}_i$. This is merely an application of Lemma 1 and Proposition 1.

Corollary 1 *Assume G_i takes the form in (4) with $f_i(0) = 0$ and $f_i'(\cdot) > 0 \geq f_i''(\cdot)$. Let $\overline{a}_i$ denote the unique solution to*

$$\frac{f_i(\overline{a}_i)}{f_i'(\overline{a}_i)} = \frac{v_i}{e}. \tag{17}$$

Then, any action no greater than $\overline{a}_i$ can be implemented by appropriately designing the assignment rule.

Proof. See the Appendix. ■

The assumption that $f_i(0) = 0$ is made only to ensure that (17) has a solution. If $\frac{f_i(0)}{f'(0)}$ is large relative to v_i, then the agent's productivity at $a_i = 0$ is already large relative to his marginal productivity and to his valuation of the prize. In this case, it is impossible to incentivize the agent to become active in the contest.

To illustrate, consider a symmetric version of Fullerton and McAfee's (1999) model with $f_i(a_i) = f(a_i)$ for all $i \in N$ and assume that $f(0) = 0$. Compare (17) to the possible outcomes of a model that assumes that the CSF takes the form in (2). In such a model, it follows from Fu and Wu (2020) that a_i can be no larger than the solution to

$$\frac{f(a_i')}{f'(a_i')} = \frac{v_i}{4}.$$

By concavity of f, $a_i' < \overline{a}_i$.

In Fullerton and McAfee's (1999) model, an unbiased contest is a special case of (2) and so it follows that actions in the Nash Equilibrium of the unbiased contest must be bounded above by a_i' for all $i \in N$. For instance, if $f(a_i) = a_i$

then $\overline{a}_i = \frac{v_i}{e}$ whereas $a_i' = \frac{v_i}{4}$. Here, $\overline{a}_i$ is 47% larger than a_i'. This provides some illustration of the potential gains from biasing the contest; it is possible to induce at least one chosen agent to work considerably harder than in the unbiased contest. The 47% is a conservative estimate since a_i' is itself an upper bound.

In fact, the set of implementable actions in Fu and Wu (2020) is a subset of $\mathcal{A}_N$. The reason is as follows. Fu and Wu (2020) show that the frontier of the set of actions that can be implemented in their setting requires zero head starts.[13] Thus, (2) reduces to

$$\widehat{p}_i(\mathbf{a}) = \frac{b_i f_i(a_i)}{\sum_{j=1}^{n} b_j f_j(a_j)}.$$

However, it can easily be verified that such a CSF can be feasibly reproduced in the Fullerton and McAfee (1999) model by giving each agent i an adjusted score of the form

$$s_i^{FW}(q_i) = \beta_i \ln G_i(q_i|a_i^*), \tag{18}$$

where a_i^* is the equilibrium action. Thus, anything that can be implemented using $\widehat{p}_i(\mathbf{a})$ can also be implemented in the current model by picking a specialized scoring function. Since the current approach allows more action profiles to be implemented, it should be expected that the optimal action profile generally differ in the two approaches.

In their setting, Fu and Wu (2020) show that a_i' is attained only if the handicaps are chosen such that agent i wins with probability $\frac{1}{2}$ in equilibrium. This occurs when $\beta_1 = \beta_2$. Thus, in contests with $n = 2$ agents and with an objective function that satisfies AIM, it is optimal to calibrate the handicaps to create a perfectly level playing field. Then, each agent wins with equal probability in equilibrium and $a_i = a_i'$ for $i = 1, 2$. Since it is impossible to force a_i higher, $B(a_1, a_2)$ is maximized. The details of B are *irrelevant* as long as it satisfies AIM. However, these result do not hold when one takes the approach proposed in the current paper.

To begin, the action profile (a_1', a_2') is also on the frontier of $\mathcal{A}_N$ when $n = 2$. This action profile is achieved when $\tau_1 = \tau_2$. In this case, the intercepts in (15) cancel out and the scores are thus ranked in the same way as (18) ranks scores

[13] Fu and Wu (2019b) and most of the prior literature restrict head starts to be non-negative. However, Drugov and Ryvkin (2017) show that negative head starts may be better.

when $\beta_1 = \beta_2$. Other action profiles on the frontier include $(\overline{a}_1, 0)$ and $(0, \overline{a}_2)$. The latter are not feasible using Fu and Wu's (2020) approach, again because there a_i is bounded above by $a'_i < \overline{a}_i$. Consider now the special case where $B(a_1, a_2) = \omega(a_1) + \omega(a_2)$, as is the case if the designer is interested in total effort or in total performance. Assume that ω is increasing and adopt the normalization that $\omega(0) = 0$. Then, $B(a'_1, a'_2) = \omega(a'_1) + \omega(a'_2)$ whereas $B(\overline{a}_1, 0) = \omega(\overline{a}_1)$. Since a'_2 becomes arbitrarily small as $v_2 \to 0$, it follows that $B(\overline{a}_1, 0)$ is unambiguously larger than $B(a'_1, a'_2)$ if v_2 is small enough. Hence, the action profile (a'_1, a'_2) cannot generically be optimal. Similarly, the optimal action profile is sensitive to the properties of ω.

Since (a'_1, a'_2) is generally not the optimal action profile, it follows that τ_1 and τ_2 are generally different. By Proposition 4, the optimal CSF is therefore not a generalized lottery CSF. Likewise, the two agents do not win with the same probability ex ante. Indeed, the optimal design and the resulting winning probabilities are sensitive to the objective function.

6 Discussion

This section discusses technical aspects of the model. First, it is noted that the distributions functions play dual roles in the stochastic performance framework. The implications of this observation are discussed. Technical extensions to the model are then considered.

6.1 Productivity and incentives

Even in the unbiased case, the primitive $G_i(q_i|a_i)$ plays a dual role. First, it describes the agent's performance or productivity, which in many applications is of direct interest to the principal. Second, from (3), it helps shape incentives for agent i and his rivals alike and is almost certainly of interest to the principal for that reason as well. For instance, in the best-shot model $f_i(a_i)$ both determines the agent's expected performance and enters the CSF in (1) directly. Thus, it is hard to disentangle productivity and incentives for comparative statics purposes.

There is a literature that examines how changes in the *sensitivity* of $f_i(a_i)$

effects the equilibrium and thereby total effort, see e.g. Nti (2004) and Wang (2010). As a comparative statics exercise this is of course interesting but, as mentioned, in many applications it is productivity and not effort that matters. Thus, it is also important to keep in mind that the *level* of $f_i(a_i)$ is likely to change when the sensitivity changes. Moreover, as $G_i(q_i|a_i)$ or $f_i(a_i)$ are primitives of the stochastic performance model, it is harder to justify thinking of changes in their sensitivity as being a design choice. See, however, the next section for one possible explanation.

It is also commonplace in the literature to compare the generalized lottery CSF to the all-pay auction CSF. In the latter, the prize is awarded to the agent with the highest action rather than the highest performance. For instance, Franke, Leininger, Wasser (2018) compare the performance of head starts and handicaps in the two settings. While this is again a legitimate comparative statics exercise, it is important to note that the designer cannot freely pick between the two under the stochastic performance premise. Again, the premise is that the action is not observable. In particular, when $G_i(q_i|a_i)$ is non-degenerate, there typically does not exist an assignment rule that produces the all-pay CSF.

6.2 Richer models of stochastic performance

It has been assumed for conceptual and technical simplicity that performances are independent. This assumption is built into existing microfoundations for the generalized lottery CSF but there is no conceptual reason to insist on this assumption more generally. Technically, however, the incentive compatibility problem becomes more complicated when correlation is permitted. This technical issue is left for future research.

Another simplifying assumption is that q_i is one-dimensional. In some applications, it may be more reasonable to assume that q_i is a vector. The analysis leading to (13) and (14) still applies, meaning that the candidate for an optimal assignment rule remains the same. To understand this, note that the likelihood-ratio is a scalar even if performance is multi-dimensional. Hence, comparing likelihood-ratios remain key. Checking incentive compatibility of the assignment rule may be more complicated, however. See Conlon (2009) and Kirkegaard

(2017) for ways of justifying the first-order approach with many-dimensional signals. If the likelihood-ratios are increasing along each performance dimension, then the condition in Lemma 2 is satisfied on an increasing set, in the sense of Conlon (2009). If distribution functions satisfy his CISP condition, then the first-order conditions are sufficient. The same holds if the likelihood-ratio has more structure but the distribution functions satisfy Kirkegaard's (2017) weaker LOCC condition. See Jung and Kim (2015) for an alternative approach that is founded more directly on the distribution of the likelihood-ratios.

Imagine that the vector of signals contains the part of the agent's performance that is of direct interest to the principal (e.g. the salesman's volume of sales) as well as an additional but not directly relevant signal (say the results of a customer satisfaction survey). The extra information cannot hurt the designer and it seems likely to be strictly beneficial.[14] Thus, there may be a link between the amount of available information and the question of the sensitivity of the contest to the agent's action that was mentioned in the previous subsection.

7 Conclusion

This paper pursues a model of contests that is based on stochastic performance. One narrow use of the model is to provide microfoundations for the popular generalized lottery CSF when the contest is unbiased. However, it does not justify the ways in which biases and design instruments are often modelled currently.

The stochastic performance premise implies that a CSF is little more than a tool that takes the uncertainty about performances and summarizes it succinctly into ex ante winning probabilities based on actions. In other words, the CSF provides a reduced-form description of the contest but it is not itself truly a primitive of the problem. The CSF can thus perhaps be useful in an intermediate step in the analysis but there is no reason to be blinded by some perceived need to specify a CSF directly or indeed to commit to one specific functional form. This mirrors the view put forward in the conclusion of Skaperdas (1996) that:

[14] Holmström's (1979) informativeness principle does not directly apply as there are no wages in the current model. However, it should be expected that the availability of more information will make it possible to implement higher actions.

> "Although helpful, axiomatizations by themselves are unlikely to settle the issue of appropriateness of a CSF for any particular contest situation."

The present paper advocates the view that the stochastic performance model is inherently compelling on its own. It need not be thought of merely as a tool whose only purpose is to justify a very particular CSF by relying on special distributions with knife-edge properties. The general model is worthy of study in its own right, all concerns about CSFs aside.

In fact, the CSF can be sidestepped entirely. Instead, the paper designs optimal contests from first principles by returning to the primitives of the problem, which, taking the microfoundations seriously, are described by stochastic performance. This yields a fundamental description of optimal assignment rules that is remarkably robust to both the designer's objective and the distribution of performances. The resulting CSF can then in principle be derived by integrating out the uncertainty over performances, but there appear little reason to take this extra step. In any case, the optimal CSF is sensitive to the distribution of performances. In a similar vein, starting from the primitives makes it possible to design optimal contests for environments that even in the unbiased case do not produce neat CSFs like the generalized lottery CSF.

In summary, a main message of the paper is that the ideal approach, whenever possible, is to start from the primitives of the contest and to build the analysis from the ground up. This simple message has implications beyond the setting considered in the current paper and several new research questions present themselves as a result. For instance, what is the optimal design of contests with multiple fixed prizes? Or of contests with multiple rounds? Such questions have been considered in the literature before but seemingly not by starting from the stochastic-performance foundation. Thus, the paper calls for more research into the full implications of contests with stochastic performance.

References

Bastani, S., Giebe, T., and O. Gürtler (2019): "A General Framework for Studying Contests," CESifo Working Paper No. 7793.

Baye, M.R. and H.C. Hoppe (2003): "The strategic equivalence of rent-seeking, innovation, and patent-race games," *Games and Economic Behavior*, 44: 217–226.

Baye, M.R., Kovenock, D., and de Vries, C.G. (1993): "Rigging the Lobbying Process: An Application of the All-Pay Auction," *The American Economic Review*, 83 (1): 289–294.

Che, Y-K. and I. Gale (2003): "Optimal Design of Research Contests," *The American Economic Review*, 93 (3): 646-671.

Chowdhury, S.M, Esteve-González, P. and A. Mukherjee (2019): "Heterogeneity, Leveling the Playing Field, and Affirmative Action in Contests," mimeo.

Clark, D. J. and C. Riis (1996): "On The Win Probability in Rent-Seeking Games," mimeo.

Clark, D. J. and C. Riis (1998): "Contest success functions: An extension," *Economic Theory*, 11: 201-204.

Conlon, J.R. (2009): "Two new Conditions Supporting the First-Order Approach to Multisignal Principal-Agent Problems," *Econometrica*, 77 (1): 249-278.

Corchón, L.C. and M. Dahm (2011): "Welfare maximizing contest success functions when the planner cannot commit," *Journal of Mathematical Economics*, 41: 309-317.

Corchón, L.C. and M. Serena (2018): "*Contest theory*," in Handbook of Game Theory and Industrial Organization, Volume II, edited by Luis C. Corchón and Marco A. Marini, Edward Elgar Publishing.

Drugov, M. and D. Ryvkin (2017): "Biased contests for symmetric players," *Games and Economic Behavior*, 103:116-144.

Epstein, G.S., Mealem, Y. and S. Nitzan (2011): "Political culture and discrimination in contests," *Journal of Public Economics* (95): 88-93.

Fain, J.R. (2009): "Affirmative Action Can Increase Effort," *Journal of Labor Research*, 30: 168–175.

Franke, J. (2012): "Affirmative action in contest games," *European Journal of Political Economy*, 28: 105-118.

Franke, J., Kanzow, C., Leininger, W. and A. Schwartz (2013): "Effort maximization in asymmetric contest games with heterogeneous contestants," *Economic Theory*, 52: 589-630.

Franke, J., Leininger, W. and C. Wasser (2018): "Optimal favoritism in all-pay auctions and lottery contests," *European Economic Review*, 104: 22–37.

Fu, A. and Z. Wu (2019): "*Contests: Theory and Topics*," in Oxford Research Encyclopedia of Economics and Finance, DOI: 10.1093/acrefore/9780190625979.013.440.

Fu, A. and Z. Wu (2020): "On the optimal Design of Biased Contests," *Theoretical Economics*, forthcoming.

Fullerton, R.L. and R.P. McAfee (1999): "Auctioning Entry into Tournaments," *Journal of Political Economy*, 107 (3): 573–605.

Grossman, S.J. and O.D. Hart (1983): "An Analysis of the Principal-Agent Problem." *Econometrica*, 51 (1): 7-45.

Jia, H. (2008): "A stochastic derivation of the ratio form of contest success functions," *Public Choice*, 135 (3/4): 125–130.

Hirschleifer, J. and J. G. Riley (1992): "*The Analytics of Uncertainty and Information*," Cambridge University Press.

Holmstrom, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*, 10 (1): 74–91.

Holmström, B. (1982): "Moral Hazard in Teams," *The Bell Journal of Economics*, 13 (2): 324–340.

Jung, J.Y. and S.K. Kim (2015): "Information space conditions for the first-order approach in agency problems," *Journal of Economic Theory*, 160: 243–279.

Kirkegaard, R. (2012): "Favoritism in Asymmetric Contests: Head Starts and Handicaps", *Games and Economic Behavior*, 76 (1): 226–248.

Kirkegaard, R. (2017): "A Unifying Approach to Incentive Compatibility in Moral Hazard Problems", *Theoretical Economics*, 12: 25-51.

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

Lazear, E. P. and S. Rosen (1981): "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89 (5): 841–864.

Mealem, Y. and Nitzan, S. (2016): "Discrimination in contests: a survey," *Review of Economic Design*, 20: 145–172.

Nti, K.O. (2004): "Maximum efforts in contests with asymmetric valuations," *European Journal of Political Economy*, 20: 1059–1066.

Rogerson, W.P. (1985): "The First-Order Approach to Principal-Agent Problems," *Econometrica*, 53 (6): 1357-1367.

Skaperdas, S. (1996): "Contest success functions," *Economic Theory*, 7: 283-290.

Tullock, G. (1980). "*Efficient Rent Seeking*," in J.M. Buchanan, R.D. Tollison and G. Tullock (eds), Towards a Theory of a Rent-Seeking Society, College Station, TX: Texas A&M University Press, 97–112.

Vojnonić, M. (2015): "*Contest Theory*," Cambridge University Press.

Wang, Z. (2010): "The Optimal Accuracy Level in Asymmetric Contests," *The B.E. Journal of Theoretical Economics*, 10 (1), Topics, Article 13.

Appendix: Omitted proofs

Proof of Lemma 1. For the first part, it has already been noted in the main text that the first-order condition at $a_i = a_i^t$ cannot be satisfied if (10) is violated. Assume next that it is satisfied. Consider a threshold rule with threshold $q_i = \underline{q}_i$. Then, the agent never wins, regardless of his performance. Hence, (7) is strictly negative at $a_i = a_i^t$. By continuity, there must then exist some threshold between $\underline{q}_i$ and $\widehat{q}_i(a_i^t)$ for which (7) is exactly zero when evaluated at $a_i = a_i^t$.

For the second part, for any action for which (10) binds, the assignment rule is essentially unique in its description of $P_i(q_i, \mathbf{q}_{-i})$ because the threshold rule maximizes (7). Thus, any assignment rule that differs on a set of performances of positive measure would fail to satisfy the agent's first-order condition. In contrast, $P_i(q_1, q_2, ..., q_n)$ is not unique when (10) is slack. The first part of this proof already identifies a threshold rule that implements a_i. By similar reasoning, there is another threshold rule with threshold above $\widehat{q}_i(a_i^t)$ that satisfies the first-order condition. An alternative rule that can be calibrated to work is to randomize between two thresholds, one for which (7) is strictly positive and one for which it is strictly negative. Such thresholds exist by the assumption that (10) is slack. Note that randomizing between two assignment rules is just a convex combination of the two and is therefore in itself an assignment rule. Thus, there are infinitely many assignment rules that satisfy the first-order condition. ■

Proof of Proposition 2. Given (13), the score is non-increasing in q_i if $\mu_i \leq 0$ for any $i \in N\backslash\{1\}$. In this case, $P_i(\mathbf{q})$ is non-increasing in q_i. By the MLRP, the agent then has an incentive to deviate downwards. This violates the incentive constraint in (11). Thus, $\mu_i > 0$ and $P_i(\mathbf{q})$ is therefore increasing in q_i. Since $\mu_1 = 1 > 0$, it also holds that $P_1(\mathbf{q})$ is increasing in q_1.

Now consider an interior action profile $\mathbf{a}$ for which $V(a_1, \mathbf{a}_{-1}) = 0$. Then, the first-order conditions of all the agents are satisfied. Since the assignment rule is monotonic, the agents' utility maximization problems are concave by the CDFC. Hence, the assignment rule is incentive compatible for all agents. By the same arguments as in Lemma 1 the assignment rule that implements the action profile is then essentially unique and given by (13).

If $V(a_1, \mathbf{a}_{-1}) < 0$ then it is impossible to design a contest that is incentive compatible for agent 1. For instance, it follows from Lemma 1 that $V(a_1, \mathbf{a}_{-1}) < 0$ if $a_1 > \overline{a}_1$, regardless of $\mathbf{a}_{-1}$. If $V(a_1, \mathbf{a}_{-1}) > 0$, then by continuity there must exist a higher a_1, closer to $\overline{a}_1$, for which $V(a_1, \mathbf{a}_{-1}) = 0$. Thus, if the action profile is on the frontier of the set of implementable actions then $V(a_1, \mathbf{a}_{-1}) = 0$ must necessarily hold. In this case, it has just been shown that the assignment rule is essentially unique. ■

Proof of Proposition 3. From (12), pointwise maximization requires that the prize is assigned to the agent with the highest value of $\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q})$, provided this is positive. The prize must be withheld if $\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q})$ is negative for all i. Is it possible that the object is not assigned if $\mu_i v_i L_i(q_i|a_i) > 0$ for some i? In this case, the feasibility constraint is slack and $\eta(\mathbf{q})$ must be zero by complementary slackness. However, if $\eta(\mathbf{q}) = 0$ then $\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q})$ is strictly positive and the prize must thus be assigned. These two statements are contradictory. Hence, the prize must be assigned whenever there is an agent with a positive value of $\mu_i v_i L_i(q_i|a_i)$.

On the other hand, imagine that $\mu_i v_i L_i(q_i|a_i) < 0$ for all i. Since the feasibility constraint is an inequality constraint the multiplier can be signed immediately, with $\eta(\mathbf{q}) \leq 0$. Thus, $\mu_i v_i L_i(q_i|a_i) + \eta(\mathbf{q}) < 0$ for all i and the prize must be withheld. In conclusion, the prize is assigned if and only if there is an agent with a positive value of $\mu_i v_i L_i(q_i|a_i)$, and in this case it is awarded to the agent with the highest score. This is the assignment rule described in (14).

If $\mu_i \leq 0$ then scores are non-increasing in q_i. Hence, $P_i(\mathbf{q})$ is non-increasing in q_i. Given the MLRP, this violates the incentive constraint. Thus, $\mu_i > 0$ for all $i \in \mathcal{N}/\{1\}$, while $\mu_1 = 1$ is positive by definition as well. Hence assignment rules are monotonic. In the case where there are inactive agents beyond the $\mathcal{N}$, or $N/\mathcal{N} \neq \emptyset$, their incentives to remain inactive are easily ensured. If the prize is not assigned to any of the active agents, then simply give it to an inactive agent at random. Such an agent can then not influence the probability that he wins the prize and his unique best response is to be inactive. There are a plethora of such rules that work. Thus, to clarify, the claim of uniqueness in the proposition refers only to the assignments of the active agents in $\mathcal{N}$. The rule in (14) is independent of the performances of the inactive agents. The same arguments as in Proposition 2 concludes the proof. ■

Proof of Proposition 4. Agent i beats agent j if performances are such that $\mu_i v_i L_i(q_i|a_i^*) > \mu_j v_j L_j(q_j|a_j^*)$, or

$$q_j < H_j^{-1}\left(e^{\frac{(\tau_i-\tau_j)}{\tau_j f_j(a_j^*)}} H_i(q_i)^{\frac{\tau_i f_i(a_i^*)}{\tau_j f_j(a_j^*)}}\right).$$

Note that the argument of the inverse function on the right is between zero and one because $\tau_i \leq \tau_j$ by assumption. Given q_i, the probability of agent i beating agent j is therefore

$$H_j(q_j)^{f_j(a_j)} < e^{\frac{(\tau_i-\tau_j)f_j(a_j)}{\tau_j f_j(a_j^*)}} H_i(q_i)^{\frac{\tau_i f_i(a_i^*)}{\tau_j f_j(a_j^*)} f_j(a_j)},$$

which obviously depends on agent j's action, a_j. Agent i must beat all rivals in order to win. Taking the expectation over q_i, his probability of winning is therefore

$$\begin{aligned}\widehat{p}_i(\mathbf{a}|\boldsymbol{\tau}) &= \int_{\underline{q}_i}^{\overline{q}_i} \prod\nolimits_{j\neq i}\left(e^{\frac{(\tau_i-\tau_j)f_j(a_j)}{\tau_j f_j(a_j^*)}} H_i(q_i)^{\frac{\tau_i f_i(a_i^*)}{\tau_j f_j(a_j^*)} f_j(a_j)}\right) f_i(a_i) H_i(q_i)^{f_i(a_i)-1} h_i(q_i) dq_i \\ &= \left(\prod\nolimits_{j\neq i} e^{\frac{(\tau_i-\tau_j)f_j(a_j)}{\tau_j f_j(a_j^*)}}\right) \int_{\underline{q}_i}^{\overline{q}_i} f_i(a_i) H_i(q_i)^{\sum_{j\neq i}\frac{\tau_i f_i(a_i^*)}{\tau_j f_j(a_j^*)} f_j(a_j)} H_i(q_i)^{f_i(a_i)-1} h_i(q_i) dq_i,\end{aligned}$$

which simplifies to (16). ■

Proof of Corollary 1. In the best-shot model, where $\widehat{q}_i(a_i^t) = H^{-1}\left(e^{-\frac{1}{f_i(a_i^t)}}\right)$ or $H(\widehat{q}_i(a_i)) = e^{-\frac{1}{f_i(a_i^t)}}$, (9) is

$$\overline{U}_i(a_i) = v_i\left(1 - e^{-\frac{f_i(a_i)}{f_i(a_i^t)}}\right) - a_i$$

and (10) simplifies to

$$\frac{f_i'(a_i^t)}{f_i(a_i^t)} \geq \frac{e}{v_i}.$$

By concavity, the left hand side is decreasing. Hence, the condition is satisfied if and only a_i^t is no greater than the solution to (17). By Proposition 1, it is then possible to implement the action. ■