Paper:
PELOSSE, Y. (2014). Dynamic Difference-Form Contests. *Journal of Public Economic Theory, 16*(3), 401-426.

http://dx.doi.org/10.1111/jpet.12067

1. Introduction

Many applications in economics involve a contest in which participants expend resources or effort to win an indivisible prize.[1] In labor tournament models, Green and Stokey (1983) and Nalebuf and Stiglitz (1983) analyze perfectly discriminatory contests where an organizer awards a prize to the contestant with the highest perceived output. This model has been extended to analyze political lobbying and research contests see e.g. Hillman and Riley (1989) and Che and Gale (1998). A second approach considers imperfectly discriminating contests in which the player with the highest observable output has a higher probability of winning, but the other contestant still has a chance to win. Dasgupta and Nti (1998), Epstein and Nitzan (2006) and Corchón and Dahm (2011), show that various probabilistic contests can be derived as an optimal choice of the contest designer. A common assumption of these two strands of literature is that all contestants choose effort levels simultaneously and independently of each other.

In reality, however, there are many situations in which the contest spans several periods, and contestants compete over a long run prior to being evaluated for awards. Examples include career promotions, lobbyists or firms competing for a contract. In this paper, we explore the optimal allocation of a prize when players compete in a finite number of periods and the winner is determined as a function of the players' stream of outputs. In each period, the participants' effort gives rise to stochastic outputs. The planner hands out an indivisible prize at the end of the game. Thereafter we call this structure a *multi-period contest*. In our multi-period contest, the organizer aggregates the observable outputs from every period and awards a prize to the winner who is determined as the outcome of a lottery with probabilities based upon the contestants's final outputs. This selection mechanism yields success functions (SFs) that map contestants' observable outputs into the likelihood of every contestant winning the prize.

We analyze the optimal SF in two complementary versions of a simple multi-period contest. The contest literature has been reviewed by Corchón (2007) and Konrad (2009). In both approaches, we assume that the administrator values the contestants' noisy output levels. This assumption is in line with the literature (see e.g., Dasgupta and Nti (1998), Epstein and Nitzan (2006), Corchón and Dahm (2010) and Gradstein (2002) and Fu and Lu (2012)) and

[1] The contest literature has been thoroughly reviewed by Corchón (2007) and Konrad (2009).

was already used by Lazear and Rosen (1981). We first study the benchmark case in which we examine which reward policy minimizes the prize implementing a maximal effort during the whole game. Our main result establishes that a version of the linear piecewise difference-form SF used by Che and Gale (2000) is the optimal design of the planner. This particular SF relates the likelihood that a contestant wins with the difference in the (aggregate) perceived outputs. Next, we examine situations where the contest organizer's goal is not necessarily to elicit maximal effort from the agents, but to maximize the expected total output. In addition the planner may also put some weight on detecting the higher-skilled agent. We show that even in this case, there exists a generalized version of the Che and Gale's two-player contest SF inducing both players to work hard.

An intuition for our results is that the crucial feature distinguishing dynamic contests from simultaneous contests is that a lagging contestant may be deterred from exerting high efforts knowing that any effort he might exert to catch up with his rival will be rendered useless. If a contestant receives a prize based on the accumulated outputs and players are ignorant of each other's inputs and outputs, then a sufficiently generous prize could stimulate contestants to put in high efforts for the whole game. However, this effect would gradually disappear as the amount of information between their respective outputs increases. In this paper, we argue that the best way to correct this effect is to use a piecewise linear difference-form SF. To put it in a nutshell, difference-form SFs are superior to other schemes because they show not only who beat whom (as any rank-order system), but by how much: these are the unique SFs that exactly convert this cardinal information into a probability measure.[2]

This paper contributes to the burgeoning literature bringing a mechanism design perspective to the study of contests (see Dasgupta and Nti (1998), Epstein and Nitzan (2006) and Corchón and Dahm (2011)). These papers provide micro-foundation for SFs as the optimal choice of a contest administrator in a simultaneous contest setting. The main motivation for their work and ours are thus similar. However, the focus of these papers is different from ours. Dasgupta and Nti (1998) consider that the organizer might not distribute the prize to the contestants. Corchón and Dahm (2011) derive sev-

[2]The idea that only differences in effort matter was initially introduced by Hirshleifer (1989). Baik (1998) was the first to use a difference-form success function.

eral popular SFs in a setting where the planner chooses the winner once the contestants have already exerted their efforts and commitment to a given contest success function is not possible. By contrast, we follow Epstein and Nitzan (2006) and assume that the planner can credibly commit to award the prize: once the contestants have expended efforts and output is realized, the organizer cannot refuse to distribute the prize according to the preannounced probabilistic rules. Our first approach - the benchmark case - analyzes the role of information in the optimal noise of a contest. There, the planner benefits from the ignorance of contestants regarding each other's outputs and the degree of noise in determining the outcome of the contest is positively related to the information contestants have about their accumulated output. Our setup may give some insight into why promotion tournaments, in which agents compete with their peers on a daily basis for a prize (the coveted promotion to the next higher echelon), appears in combination with other forms of payment: in our first setup we identify an environment where compensation functions take the form of a constant prize such as a promotion or enhanced status in addition to a probabilistic scheme based on absolute differences in perceived outputs.

Our paper is also linked to the literature on the optimal prize allocation in contests. This literature is geared towards the design of elimination tournaments, in which participants are successively eliminated and the organizer can choose the optimal number of stages to elicit maximum aggregate effort. Fu and Lu (2009) show the optimality to award the entire prize purse (only) to the contestant who wins the first prize in the finale, regardless of the sequence of the contest. Here we assume that the multi-period context structure is exogenously given and cannot be adjusted by the planner. Thus, in our setting the contest organizer is not endowed with the flexibility of creating several prizes, breaking the game into a series of battles of smaller length. However, in contrast to the present model, in these papers the SF is exogenously given. Gradstein and Konrad (1999) show that the optimality of a contest design depends on how discriminatory the ratio-form contest SF is. Their analysis is thus in line with our main theme: in contests that last several stages a better outcome is always engendered for the organizer through rendering the contest sufficiently noisy.

Recently, the factors that tend to countervail the discouragement effect. in dynamic contests have also been analyzed by Konrad and Kovenock (2009, 2010). They show how the agents' stochastic ability (2010) and the use of intermediate prizes (2009) can restore a stiff competition between agents in a multistage race in which players compete in a sequence of simultaneous move component contests. In these models there is a (possibly infinite) sequence of battles between two contestants who accumulate stage victories, and the contestant who first accumulates a sufficiently larger number of such victories is awarded the prize. In another dynamic setting, Münster (2009) explores the welfare enhancing effect of a repeated contest with asymmetric information. He shows that contestants with a high ability sometimes put in little effort in an early round in order to make the opponents believe that their ability is low. Thus, a main difference with the present model is that the organizer is not allowed to spur competition by choosing the SF. Our two approaches are more closely related to Dubey and Wu (2001) and Dubey and Haimanko (2003). These authors analyze a dynamic contest model with sampling. Their focus is on the design of a random spot check device that triggers competition through maintaining some uncertainty on the identity of the winner up to the very end. To do so, they inject some noise via a (secret) sampling on the stream of outputs produced over time periods. A distinct feature of our analysis is that we rule out situations in which the organizer may ignore some of the periods of the competition. In this sense, we complement these papers to situations in which the administrator seeks to have a well-rounded, fully accurate picture of the agents' performance history.

This paper is also related to the principal-multi-agent settings when there are positive spillovers in production between agents and the principal optimally design incentives (see e.g. Lockwood (2000) and Itoh (1991)). From this perspective, the current paper provides an intriguing connection to this literature: our model mirrors a standard moral hazard setting in which the designer has the additional latitude to render the contestants technology more or less separable via the choice of a reward policy. This point is discussed in more details in section 4.1.

This paper is organized as follows. In sections 2 we give the description of the game. Section 3 presents the existence result and the characterization of implementable mechanisms. In Section 4 we completely characterize the unique optimal policy. In Section 5 we discuss the robustness of our results and conclude. All proofs are relegated in the Appendix.

2. A multi-period contest

Consider two agents 1 and 2. The agents take part in a multi-period contest which is comprised of a finite sequence of $T \geq 1$ simultaneous move games. At the start of any period $t = 1, ..., T$ agents 1 and 2 simultaneously choose an effort level e_1 (resp. e_2) in a finite set E. Let e^* be the **maximal element** of E i.e. e^* represents a player's maximal effort level. Note that the finiteness requirement of E follows from the fact that in a multi-period contest with infinite sets of effort levels, the planner may fail to implement maximal effort levels as a pure strategy equilibrium. This is a direct implication of the analysis of Dubey and Wu (2001). Consequently, we restrict our analysis to finite sets of efforts. [3]

The contestants' effort levels lead to random outputs q_1 (resp. q_2) each taking their values in a closed interval $\left[\underline{q}_i, \overline{q}_i\right] \equiv Q_i$ of $\mathbb{R}_+$. To avoid trivial situations, we confine attention to the case where the sets of outputs producible by contestants are not too different from each other and postulate that $Q_1 \cap Q_2 \neq \emptyset$. The assumption that the sets of feasible effort are symmetric across all agents and that there are a continuum of outputs is for notational simplicity only.

For each contestant i, there is a positive correlation between his effort e and his productivity q_i captured by continuous full support conditional cdfs, $G_i(\cdot \mid e)$, for all $e \in E$. More specifically, our key assumption on the cdfs of $i(= 1, 2)$, $G_i(\cdot \mid e)$, states that when an agent exerts maximal effort, $e^* \neq e$, his outputs go up, in the sense of first-order stochastic dominance i.e. $G_i(q_i \mid e^*) \leq G_i(q_i \mid e)$ for all q_i and strictly less than for some q_i.

For each period $t = 1, ..., T$, agents simultaneously choose their effort level $e_1(t)$ (resp. $e_2(t)$) with knowledge of the full history of their *own* effort choices and *own* production. However, each participant may have some arbitrary information regarding his rival's stream of outputs. Formally, let $q_i(t)$ be the agent i's output realization at period t. We define i's history at time t by $(e_i(1), ..., e_i(t-1), q_i(1), ..., q_i(t-1)) \equiv (\mathbf{e}_i(t-1), \mathbf{q}_i(t-1))$. Accordingly, we denote the history of the game at time $t = 1, ..., T$ by $h^t = ((\mathbf{e}_1(t-1), \mathbf{q}_1(t-1)), (\mathbf{e}_2(t-1), \mathbf{q}_2(t-1)))$. We assume throughout that *any contestant i can observe only some subset of the outputs produced by his rival and not the rival's inputs of effort*. For any history, h^t, let $\mathcal{I}_i(h^t)$ denote i's information set con-

[3] For example, if players had access to an infinite compact set $E \subset \mathbb{R}_+$ of effort levels, then a contestant could depart from the maximal effort level and pick some $e \neq e^*$ slightly lower than e^*, at some periods, while hardly affecting his winning probability. As a consequence, the prize inducing players to work hard would be prohibitive for the organizer.

taining h^t. Hence, this set has the property that any two histories, h^t and h'^t, have their respective i's histories, $(\mathbf{e}_i(t-1), \mathbf{q}_i(t-1))$ and $(\mathbf{e}_i^{'}(t-1), \mathbf{q}_i^{'}(t-1))$, such that $(\mathbf{e}_i(t-1), \mathbf{q}_1(t-1), \mathbf{q}_2(t-1)) = (\mathbf{e}_i^{'}(t-1), \mathbf{q}_1^{'}(t-1), \mathbf{q}_2^{'}(t-1))$ only if $h'^t \in \mathcal{I}_i(h^t)$. This captures the fact that each contestant is **ignorant** of his rival's effort history. However, notice that contestants may observe some part or the full raw of others' outputs.

Thereafter, we define $\mathscr{I}_i$ as the set of all possible information sets for agent i with this property. Given these constructions, we define a pure strategy σ_i for agent i as a mapping $\sigma_i : \mathscr{I}_i \to E$, that specifies for every information set an element of the feasible effort levels E. Let σ_i^* be the **maximal effort strategy for agent** i when $\sigma_i^*(\mathcal{I}_i(h^t)) = e^*$ for all $\mathcal{I}_i(h^t) \in \mathscr{I}_i$. This completes the description of the extensive form of the game, but it still remains to specify agents' payoffs at each terminal history.

2.1 The administrator's policy

The performance review is a continuous process. The administrator will therefore consider the *entire stream of outputs*, $(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$, obtained by agents 1 and 2 at a terminal history. Formally, a **rule**, R, is a pair of real mappings, $(R_1, R_2) \equiv R$, where each $(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) \mapsto R_i((\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$ tells the administrator how much apart is agent i from agent $-i$ at any terminal history. Because we are restricting our attention to *anonymous* (or symmetric) rules, it follows that mapping R_i is invariant under any relabeling of the agent's name.[4] Define $\mathcal{Q}_R$ as the (possibly infinite) set of all possible **leads**, $r = R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$, that can be attained at the *terminal* histories by *all* agents under $R \equiv (R_i, R_{-i})$. For instance, the most natural choice for a rule is the **additive rule**, $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) := \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$. In this special case, $\mathcal{Q}_R = [-T\overline{r}, T\overline{r}]$ with $\overline{r} \equiv \max\left\{\overline{q}_1 - \underline{q}_2, \overline{q}_2 - \underline{q}_1\right\}$.[5]

Conditionally on $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$, the administrator awards the prize $B \circ R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$ with probability $F \circ R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$ for $(i = 1, 2)$.[6]

Accordingly, we define a **reward policy** (or simply a policy) as a tuple of mappings, $R \equiv (R_i, R_{-i})$, $F_R : \mathcal{Q}_R \to [0, 1]$ and $B_R : \mathcal{Q}_R \to \mathbb{R}_+$. B_R determines the size of the prize $B_R(r)$ (resp. $B_R(-r)$, for agent 2) and $F_R(r)$ the probability of winning this prize for agent 1 (resp. $F_R(-r)$, for agent 2) as a function of the lead r (resp. $-r$, for agent 2). Thereafter, the mapping F_R is

[4] That is we have for a bijection $\varphi : \{1, 2\} \to \{1, 2\}$, such that $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) = R_{-i}(\mathbf{q}_{\varphi(i)}(T), \mathbf{q}_{\varphi(-i)}(T))$.

[5] Recall that by assumption $\overline{q}_i > \underline{q}_{-i}$ is satisfied for $i (= 1, 2)$.

[6] Given two functions f and g, we define $f \circ g$ as the composite of f and g.

termed as the **success function** (SF, in short).
Little restriction is placed on the set of possible policies. We shall only consider mappings awarding the prize with *certainty* to *one and only one* of the agent[7] and policies that are *unconditional* on the agents' identity. Given a rule R, the set of SFs is denoted by $\mathscr{F}_R$. Formally,

$$\mathscr{F}_R := \left\{ \begin{array}{c} F_R : \mathcal{Q}_R \to [0,1] \\ F_R(r) + F_R(-r) = 1, \forall r \in \mathcal{Q}_R \end{array} \right\}.$$

Hence, generally, the probability to award the prize to an agent i is *not* perfectly correlated on the absolute lead he has achieved during the competition. We assume that each agent has the same utility function u. Each agent derives $u(B)$ units of positive utility from winning the prize B. Thus, given a reward policy (F_R, B_R), there is a function $F_R(r)u_i(B_R(r)) \equiv u_{i|F_R} \circ B_R(r)$, which gives the expected utility when i reaches a lead $r \in \mathcal{Q}_R$ at a terminal history. Let us point out that in this paper we assume that the reward –potentially based on the margin of victory – takes the form of a prize. Note that this means that the prize B might potentially be huge (but finite), compensating a contestant for his extremely low probability of winning at some information sets. In fact, any prize-based incentive scheme has this disadvantage. For example, if agents have disparate abilities so that only one believes he has reasonable chances to win, then the planner will have to disburse a sizable pot of money to generate competition. The present model is thus most plausible when contestants are drawn from the same population so that each player believes he has reasonably as good a chance as anyone else to win if he works.
Agents also incur a disutility from effort described by a function $d_i : (E)^T \to \mathbb{R}_+$ for $i(= 1, 2)$. We assume that maximal effort incurs the most disutility, i.e. $d_i(\mathbf{e}^*{}_i(T)) > d_i(\mathbf{e}_i(T))$ for all $\mathbf{e}_i(T) \in (E)^T \setminus \{\mathbf{e}_i^*(T)\}$. Accordingly, agent i's overall utility equals $u(B) - d_i(\mathbf{e}_i(T))$ if he wins and expends a stream of effort levels $\mathbf{e}_i(T)$ over the T periods, and $u(0) - d_i(\mathbf{e}_i(T))$ otherwise with $u(0) = 0$. We further postulate that agents value the prize sufficiently so that $u(B) \to \infty$ when $B \to \infty$.

2.2 The (reduced) normal form game

As mentioned above, in this first model, none of the two agents is informed about his rivals' effort history when making his own decision at any time period t. In this context, we analyze the normal (or strategic) form game that

[7]This appears to be necessary if one wants to guarantee the credibility of such a mechanism.

arises from the above extensive form game.[8] Hence, we identify the equivalent pairs of (unreduced) pure strategy profiles as follows. We may consider two pure strategies of agent i, σ_i and $\sigma_i^{'}$ as equivalent if the set of histories h^t's that cannot be reached under σ_i and $\sigma_i^{'}$ is equal and if these two strategies agree at all the other information sets i.e. $\sigma_i(\mathcal{I}_i(h^t)) = \sigma_i^{'}(\mathcal{I}_i(h^t))$ whenever h^t is reachable under σ_i and $\sigma_i^{'}$. Thereafter we will denote the set of such reduced strategies for any agent i as $\sum$ and the set of all reduced strategy profiles as $\overline{\sum}$. Each pure unreduced strategy profile $\sigma = (\sigma_1, \sigma_2)$ and policy from the organizer induces a probability distribution on the set of terminal histories i.e. the set of all histories h^{T+1}.[9] Since unreduced equivalent strategies must reach and agree on the same histories, this implies that equivalent pairs induce the same expected payoff.

Denote the probability for agent i to win the prize B under profile σ under a reward policy (F_R, B_R) by $p_i(\text{i wins} \,|\, \sigma; F, B)$ and i's expected disutility by $D_i(\sigma_i)$.[10] Given the strategy profile σ, $p_i(\text{i wins} \,|\, \sigma; F_R, B_R)u(B) - D_i(\sigma_i) \equiv \Pi_i(\sigma)$ is agent i's expected payoff when the reward policy is (F_R, B_R). Since we hold the time horizon of the contest fixed throughout, $\Gamma(\mathscr{I}, (F_R, B_R))$ denotes the *reduced* normal form game induced by a reward policy (F_R, B_R) under the set of information partitions $\{\mathscr{I}_i\}_{i\in\{1,2\}} \equiv \mathscr{I}$. Thus, each $\Gamma(\mathscr{I}, (F_R, B_R))$ is a simultaneous-move game and the solution concept we use up to section 4 is that of a **pure-strategy Nash equilibrium (NE)** of this game.[11] It will prove convenient to use the notation $\sigma \setminus \tilde{\sigma}_i \equiv (\tilde{\sigma}_i, \sigma_{-i})$. Hence, when the organizer chooses a policy (F_R, B_R), we say that σ is a Nash equilibrium (NE) of $\Gamma(\mathscr{I}, (F_R, B_R))$ if

$$\Pi_i(\sigma) \geq \Pi_i(\sigma \setminus \tilde{\sigma}_i) \text{ for any } \tilde{\sigma}_i \in \sum \text{ and } i = 1, 2.$$

In the sequel we say that a reward policy (F_R, B_R) is **NE-feasible** if it implements σ^* as a NE of $\Gamma(\mathscr{I}, (F_R, B_R))$ and if the size of the bonus B is *finite*. Define $B_R(F_R) \equiv \min\{B : \sigma^* \text{ is a NE of } \Gamma((\mathscr{I}), (F_R, B_R))\}$.

3. Existence of a NE-feasible reward policy

In the present model, the information of a participant about the rival pertains

[8] Here, we follow the methodology used in Dubey and Wu (2001).

[9] For the sake of notational consistency, we regard $H^1 = h^1$ as the "dummy" singleton history for period 1.

[10] When i is ignorant of his rival, it is easy to check that i's expected disutility is independent of σ_{-i}.

[11] In Section 5 we analyze multi-period contests modeled as perfect information trees which we analyze for its subgame perfect equilibria in behavioral strategies.

only to the outputs produced. The key problem for the contest administrator is that, when contestants can observe their rivals' outputs over a sufficiently large number of periods, it is not possible to implement σ^* a NE because there exist some histories where they win (resp. lose) without having to expend any further effort. In the following we characterize the set of all reward policies that implement σ^* as a NE of $\Gamma(\mathscr{I}, (F_R, B_R))$ with finite prizes and under various information settings.

To this end, define $I_r \subset [-\overline{r}T, \overline{r}T]$ as the interval centered at r of length $2\overline{r}$ with $\overline{r} \equiv \min\left\{\overline{q}_1 - \underline{q}_2, \overline{q}_2 - \underline{q}_1\right\}$ and consider,

$$\{I_r : r = -k^*(\mathscr{I})\overline{r}, ... - 4\overline{r}, -2\overline{r}, 0, 2\overline{r}, 4\overline{r}, ..., k^*(\mathscr{I})\overline{r}\} \equiv \mathscr{P}^*(\mathscr{I}),$$

with $T - N^*(\mathscr{I}) \equiv k^*(\mathscr{I})$. Here, $N^*(\mathscr{I})$ denotes the largest number of periods for which both participants have always some incentives to exert maximal effort e^* in the game $\Gamma(\mathscr{I}, (F_R, B_R))$ when F_R is of the all-pay auction type.[12] Finally, recall that $u_{i|F_R} \circ B_R(r)$ is the expected utility when i reaches a lead r at a terminal history under (F_R, B_R).

Given an information partition, $\mathscr{I}$, one may define the class of rules to evaluate the difference in the aggregate sum of agents' output histories. Consider the (linear) order, $\succeq_{i,\epsilon}^{\mathscr{I}}$, on $Q_1^T \times Q_2^T$ for all $i (= 1, 2)$ and define,

$$\begin{aligned} &(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) \succeq_{i,\epsilon}^{\mathscr{I}} (\mathbf{q}_i^{'}(T), \mathbf{q}_{-i}^{'}(T)) \\ \text{iff } &\sum_{\tau=1}^{k^*(\mathscr{I})} q_i(\tau) - q_{-i}(\tau) \geq \sum_{\tau=1}^{k^*(\mathscr{I})} q_i^{'}(\tau) - q_{-i}^{'}(\tau) + \epsilon, \end{aligned}$$

for some $\epsilon \in (0, \overline{r}]$. Consider the class of functions, $R_i : Q_1^T \times Q_2^T \to \mathbb{R}$ that represents[13] $\succeq_{i,\epsilon}^{\mathscr{I}}$ for all i. We will say that the pair of functions, R, is ϵ-**additive** if every R_i represents $\succeq_{i,\epsilon}^{\mathscr{I}}$ on $\mathcal{Q}_R$. When R is an ϵ-additive rule, we call (F_R, B_R) a ϵ-**difference-form** policy (or **difference-form** policy for short). The following theorem characterizes the set of NE-feasible policies for any information partition, $\mathscr{I}$.

Theorem 1 *Consider a $T \geq 2$ multi-period contest game $\Gamma(\mathscr{I}, (F_R, B_R))$. A reward policy (F_R, B_R) is NE-feasible for $\Gamma(\mathscr{I}, (F_R, B_R))$ if and only if $(F_R, B_R) \in \mathscr{F}_R \times \mathscr{B}_R$ is a difference-form policy such that $u_{i|F_R} \circ B_R : \mathcal{Q}_R \to \mathbb{R}_{++}$ is (non constant) increasing over each $I_r \in \mathscr{P}^*(\mathscr{I})$.*

[12]The formula for $N^*(\mathscr{I})$ is given below.

[13]In other words, R_i is an order-preserving function.

Proof. See Appendix A.■

Theorem 1 says that as long as the rule is ϵ-additive and the set of discontinuity points is judiciously distributed over $\mathcal{Q}_R$, we are guaranteed that both contestants affect their probability of winning by exerting maximal effort at each of their information set. This result is also important since it shows that difference-form policies are the only possible optimal policies. In other words, the administrator is always best when he considers policies that build only on (possibly very rough) estimate of differences in the sum of accumulated output across all T periods.

A direct consequence of Theorem 1 is that in our study of the optimal prize(s) and success function(s), it will suffice to consider difference-form policies under the simple rule, $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) = \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$. This follows since NE-feasible policies are defined modulo any ϵ-additive rules. Accordingly, we will simply denote a difference-form policy by (F, B).

Moreover, an important case of Theorem 1 is when the prize function B is a *constant function*. In this case, as an immediate corollary of Theorem 1, we see that the set of NE-feasible reward policies $\mathscr{F}^{NE}(\mathscr{I})$ (modulo ϵ-additive rules) is described by

$$\mathscr{F}^{NE}(\mathscr{I}) \equiv \left\{ \begin{array}{c} F : [-\overline{r}T, \overline{r}T] \rightarrow [0,1] \\ F(r) + F(-r) = 1, \forall r \in [-\overline{r}T, \overline{r}T] \\ \text{with } F \text{ (non constant) increasing over each } I_r \in \mathscr{P}^*(\mathscr{I}) \end{array} \right\}$$

Theorem 1 reveals that SFs increasing on a sufficiently fine grid of discontinuity points are necessary to obtain a positive contest effort whenever some agents are lagging far behind. These difference-form SFs avoid contests becoming trivial because they maintain a sufficient correlation between effort and performance.

An intuition for Theorem 1 is provided by an answer to the following question.

Given an information partition $\mathscr{I}$, what about the largest number of periods which still leaves some incentive for agent i to exert maximal effort e^* at any information sets?

For the sake of exposition, the following discussion assumes policies (F_R, B_R) defined from the special additive rule, $R_i(\mathbf{q}_i(T+1), \mathbf{q}_{-i}(T+1)) := \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$. Hereafter, we simply denote a policy defined on this particular additive rule by (F, B).

Suppose contestant i can perfectly observe his rival's output during $0 \leq n_i \leq T$ periods of the contest. Let $\mathscr{I}_{n_i}$ denote a set of information sets wherein agent i perfectly observes his rival's output during n_i periods. This means, $\max\{n_1, n_2\} \equiv n^*$ is the maximal number of periods wherein a contestant can perfectly observe his rival's output. As in Dubey and Wu (2001), it is useful to define

$$\Delta^O(\mathscr{I}_{n_i}, t) = \begin{cases} \underline{q}_i(t) - \overline{q}_{-i}(t) \text{ if } \ i \text{ observes } -i\text{'s output at } t; \\ \underline{q}_i(t) - \underline{q}_{-i}(t) \text{ otherwise,} \end{cases}$$

as the **most optimistic per period lead** for contestant i over $-i$ under information partition $\mathscr{I}_{n_i}$. Let $N_i(k)$ be the set of periods wherein contestant i can observe his rival's output over the first k-periods of the game. Thus, $\sum_{\tau \in N_i(k)} \Delta^0(\mathscr{I}_{n_i}, \tau)) \equiv \overline{h}(k, \mathscr{I}_{n_i})$ is the **most optimistic lead** that i may have achieved under $\mathscr{I}_{n_i}$ at the end of period k when he has observed the *worst* possible outcome over some subset of the observable periods $N_i(k)$ i.e. i has obtained an output $\underline{q}_i(\tau)$ at periods in $N_i(k)$ but assumes his rival produced $\underline{q}_i(\tau)$ at all the unobservable periods. Denote integers by $\mathbb{Z}$ and define

$$\max\left\{k \in \mathbb{Z} : (\overline{q}_i - \underline{q}_{-i})(T - k) + \overline{h}(k, \mathscr{I}_{n_i}) \geq 0\right\} \equiv \overline{N}(\mathscr{I}_{n_i})$$

as the *largest number of periods* which still leaves some incentives for participant i to exert maximal effort e^* under information partition $\mathscr{I}_{n_i}$ when the success function is of the *all-pay auction* type and i assumes the most optimistic scenario. Arguing symmetrically, one can define **the most pessimistic per period lead** for contestant i over $-i$ when under information partition $\mathscr{I}_{n_i}$, $\Delta^P(\mathscr{I}_{n_i}, t)$, and $\underline{h}(k, \mathscr{I}_{n_i})$. In this case, the integer,

$$\max\left\{k \in \mathbb{Z} : (\underline{q}_i - \overline{q}_{-i})(T - k) + \underline{h}(k, \mathscr{I}_{n_i}) < 0\right\} \equiv \underline{N}(\mathscr{I}_{n_i}),$$

is the *largest number of periods* which still leaves some incentives for participant i to exert maximal effort e^* under his information partition $\mathscr{I}_{n_i}$ when the success function is of the all-pay auction type and i assumes the most pessimistic scenario.

Now, since the administrator seeks to avert both scenarios, we define

$$\min\left\{\underline{N}(\mathscr{I}_{n_i}), \overline{N}(\mathscr{I}_{n_i})\right\} \equiv N^*(\mathscr{I}_{n_i}).$$

Hence, given the information partition $\mathscr{I}$ the integer

$$N^*(\mathscr{I}) = \min_{i \in \{1,2\}} \left\{N^*(\mathscr{I}_{n_i})\right\}$$

is the **largest number of periods** for which both participants have always some incentives to exert maximal effort e^* in the game $\Gamma(\mathscr{I}, (F, B))$ when F is of the all-pay auction type. This means that *NE-policies can be sustained under the all-pay auction during* $N^*(\mathscr{I})$ *periods*. Hence, given the information structure – the partition $\mathscr{I}$ – there exists some scenarios – histories – in which both contestants would shirk during $T - N^*(\mathscr{I}) \equiv k^*(\mathscr{I})$ periods. To avoid these scenarios, the administrator needs to consider functions (F, B) that are non constant increasing on a sufficiently large number of intervals of $[-\overline{r}T, \overline{r}T]$.

4. Existence and characterization of the optimal reward policy

In this section, we characterize the unique NE- feasible reward policy (modulo ϵ-additive rules) which implements σ^* at the least cost for the administrator. An **optimal reward policy**, $(F^*, B(F^*))$, is a NE-feasible policy such that there is no other NE-feasible policy, $(F^{'}, B(F^{'}))$ with $B(F^{'}) < B(F^*)$. In lieu of potentially complex reward policies, we first establish that optimal reward policies will necessarily lie in the set of NE-feasible policies, $(F, \overline{B})$ where $\overline{B}$ is a *constant* function.

Lemma 1 *The pair* (F^*, B^*) *is an optimal reward policy only if* B^* *is a constant function i.e.* $B^*(r) = \overline{B}$ *for all* $r \in [-rT, rT]$.

Proof. See Appendix B ■

This result shows that even if the designer is able to assign different values to the prize (based upon the final performance of the contestants), he will always prefer a noisy allocation scheme. The intuitive explanation of this result is that by putting up a bunch of noise in the allocation of the prize – via the SF –, the contest organizer can provide contestants' incentives to supply effort for free. In effect, an increasing prize function will provide additional incentives for effort. This benefit, however, comes at a cost of having to put more money on the table. In contrast, by influencing the contestant's likelihood of winning through an imperfectly discriminatory contest, the organizer can spur agents to work hard, simply because each contestant must exert maximal effort if he wants to continue to win the prize with good probability. The upshot is that pure difference-form reward policies beat mixed reward policies i.e., policies made of a SF and a prize that are based on the margin.

We now turn to the issue of the optimal characterization of a (difference-form) policy (modulo ϵ-additive rules) (F_R^*, B_R^*) i.e. the optimal design of the

SF(s), F_R when R is the additive rule, $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) = \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$ and B^* a constant function. Thus, we consider SF(s) defined on $[-\overline{r}T, \overline{r}T]$ and simply use the notation (F, B) instead of (F_R, B_R) to denote a typical element of this class of policies.

In the following, we outline several tools used in our analysis of optimal SFs: jump of discontinuities and point of discontinuities of a success function.

A success function F is said to have a **jump discontinuity** or **jump** at z if there are leads $r_0 < z < r_1$ such that

(i) F is defined and continuous on $[r_0, z)$ and on $(z, r_1]$;
(ii) the left and right limit at z exist but are different.

Thereafter we refer to z as a **point of discontinuity** of F and for any $\epsilon > 0$, $F(z+\epsilon) - F(z-\epsilon) \equiv \mu_F(z)$ is the **jump** of F at z. There are three questions that are relevant concerning the optimal success function:

(i) How the form of the SF influences the agents incentive to exert e^* rather than $e \neq e^*$ at a given history?
(ii) What is the characterization of the set of discontinuity points of the optimal success function?
(iii) What is the characterization of the jump discontinuity of optimal SF(s)?

We discuss these three issues in turn. The answers to (i)-(iii) build on the following key observation.

Lemma 2 *The optimal SF F^* must increase at a constant rate.*

Proof. See Appendix B ■

Now let us return to (i)-(iii). Conditionally on each lead, the agents' effort choice in each period balances its marginal disutility with the expected marginal increment in the probability of winning. The marginal increment in the probability of winning depends on the *rate of increase of the success function*. Therefore, the administrator must maximize the expected marginal increment in the probability of winning at each lead. By definition, under σ^*, the agents are not allowed to incur higher marginal disutility in the first periods in order to preempt the leading position. Moreover, the set of outputs is an interval of $\mathbb{R}_+$, which entails an *infinite* number of leads that can be reached with a positive probability at the last period.

We now turn to the issue of the relation between the information agents have regarding each others and the set of discontinuity points of F^*. As noted earlier, the form of the NE-feasible SFs depends on the information partition $\mathscr{I}$. As information gets coarser between agents, the set of leads that can be reached with a positive probability and where agents get discouraged decreases. Moreover, the marginal increment in the probability of winning is increasing in the rate of increase of the SF. Hence, it is optimal to *minimize* the number of jumps. This accounts for the fact that F^* must be increasing only on the set of leads that can be perfectly observed by agents under $\mathscr{I}$. From these key properties, we deduce that, the only optimal options for the administrator is that a contestant's probability of winning the prize increases at a *constant* rate at any possible lead that might be reached at the last period. It follows from all these observations that F^* is uniquely characterized by the linear and continuous difference-form SF à la Che and Gale. We are now in a position to state the advertised result:

Theorem 2 *Consider a $T \geq 2$ multi-period contest game $\Gamma(\mathscr{I}, (F^*, B(F^*)))$. Then, the optimal success function, is the (unique) continuous piecewise linear difference-form success function,*

$$F^*(r) = \max\left\{\min\left\{\frac{1}{2} + \frac{r}{2\overline{r}k^*(\mathscr{I})}, 1\right\}, 0\right\} \text{ if } r \in [-\overline{r}T, \overline{r}T].$$

Proof. See Appendix B ■

This theorem gives a justification for the use of the family of SFs which are a generalized version of piecewise linear difference-form success function introduced initially by Che and Gale (2000). [14] Intuitively, the all-pay auction SF is not optimal in multi-period settings since it deters agents from exerting effort over time: because of its purely ordinal criterion, there exists some scenario in which contestants' winning probabilies are left invariant with respect to their output levels. This causes players to give up. With respect to other forms of SFs, win probabilities converting directly outputs into probability of winning according to an absolute criterion offer the advantage to become sensitive to the absolute effort levels – via the aggregate difference in perceived outputs.[15] From the viewpoint of the designer, the best SF is thus the

[14] It is worth mentioning that this family of SFs has been given microeconomic underpinnings by Corchón and Dahm (2010). Here, our results complement their work and show – in a very different setup – how this function arises naturally in multi-period contests.

[15] Recall that there is a positive correlation between a contestant's effort and his output under first-order stochastic dominance.

one that is the most sensitive to effort over time which, by construction, is the linear piecewise difference-form SF. In short, a piecewise difference-form SF allows the designer to make contestants *"run to keep at the same place"* at the least cost. Several other comments are in order.

First, note that the scope of our analysis has a natural limitation: we provide the optimal SF but are unable to write an explicit formula for the optimal prize. This is the price we pay for our generality e.g. we assume nothing about utility other than the fact that it is monotonically increasing in the size of the prize. Second, when we consider simultaneous contests, i.e. $T = 1$, a dramatic change occurs. In this case, it is straightforward to show that the optimal SF coincides with the all-pay auction i.e. the organizer allocates the prize with certainty to the contestant with the largest output. Intuitively, this shift between simultaneous and multi-period contests arises because in the static case the designer need not further worry about the sensitivity of the contestants to the information regarding each other's outputs. As a result, the planner can employ an all-pay auction SF to make a contestant's output infinitely sensitive at the margin, thereby eliciting maximal effort from the participants (in equilibrium) at the least cost.

An interesting question in the optimal design of contests is how the organizer determines the "noise" factors, such as the imperfections in performance measurement and evaluation. Theorem 2 suggests an answer. In the functional form derived in Theorem 2 there are two parameters in the model, $k^*(\mathscr{I})$ and $\overline{r}$, which render the allocation of the prize noisy. By contrast, the literature (see e.g. Che and Gale (2000)) analyzes a class of difference-form contests with an exogenous scalar that specifies how deterministic the contest is. The first parameter, $k^*(\mathscr{I})$, can be interpreted as an "information" parameter, with a greater value of $k^*(\mathscr{I})$ inducing less noise. Below – see section 4.2 – we shall formally examine how the optimal level of noise does indeed depend on the information players have about each other. The second parameter, $\overline{r}$, can be construed as a "technology" parameter. Intuitively, this scalar must be seen as a measure of the difference in the agents' technology. In particular, when the sets of producible outputs are concentrated in a small neighborhood of the positive numbers, the organizer tends to allocate the prize to the contestant who produces the largest performance, even if he beats his rival by a single unit of output. Intuitively, it would be harder to sustain maximal effort as a NE with a noisy ranking because this would not distinguish the slim difference in outputs between the two contestants. Conversely, when all the players have sufficiently dispersed sets of production, Q_i, i.e., $\bar{r}$ is very large, it is always better to allocate the prize according to a scheme that approaches a pure lottery. This is so because one of the two contestants may obtain a significantly larger output than his rival. Hence, awarding the prize with nearly perfectly discriminating SFs would reveal the huge lead, thereby deterring players from putting in effort.

In the literature, another popular family of success functions – the so-called ratio-form introduced by Tullock (1980) – assumes that a contestant's probability of winning the contest equals the ratio between this contestant's own perceived output and the sum of perceived outputs, or a variant of this.[16] Tullock's formulation is thus based on a *relative* criterion. In light of Theorem 1, ratio form SFs would also implement σ^* as a NE. However, difference-form SFs outperform this family because they convey a (cardinal) information about the difference in absolute perceived outputs – which is a proxy for the difference in contestants' absolute effort levels. By construction, this information is partially lost in the case of SFs based on relative outputs: unlike ratio-form SFs, piecewise linear difference-form SFs allows the designer to reward the marginal increases (resp. reductions) in the lead achieved by the contestants. The upshot is then that ratio-form SFs require a higher prize to elicit maximal effort than difference-form SFs.

4.1 Discussion

As mentioned above, our results are also related to the voluminous principal-multi-agent literature a la Mookherjee (1984) and Lockwood (2000). For instance, Lockwood (2000) shows that the optimal payment to any agent will also depend (negatively) on the aggregate output of the other agents. Similar results have been obtained in different frameworks (see e.g. Itoh (1991), Demski and Sappington (1988)). The need of an additive rule R i.e. a difference-form policy, permits to uncover the link between our setup and the above models. Indeed, it suffices to regard ϵ-additive rules as the mirror of some underlying production externalities between the two contestants: under the disguise of these rules, a "spillover" effect comes into play, enabling each contestant to increase his own "score" while reducing that of his rival.

4.2 Coarser information partitions and optimal success function

[16] For example, Alcade and Dahm (2007) introduce a family of SFs that relies both on relative efforts and absolute effort differences.

We now study the effect of coarsening the information contestants have about each others' outputs on the optimal reward policy. Recall that $\mathscr{I}_{n_i}$ denote the set of information sets wherein agent i perfectly observes his rival's output during n_i periods.

Definition 1 *Fix two information partitions, $\{\mathscr{I}_{n_1}, \mathscr{I}_{n_2}\} = \mathscr{I}$ (resp. $\left\{\mathscr{I}_{n_1'}, \mathscr{I}_{n_2'}\right\} = \mathscr{I}'$), with horizon T. Let $\min\{n_1, n_2\} = n^*$ (resp. $\min\left\{n_1', n_2'\right\} = n'^*$) be the maximal number of periods for which contestants can perfectly observe their rival's output in the game $\Gamma(\mathscr{I}, (F, B))$ (resp. $\Gamma(\mathscr{I}', (F, B))$). We say that information partition $\mathscr{I}'$ is* ***coarser*** *than $\mathscr{I}$ if $n^* > n'^*$.*

The following theorem establishes the convergence result.

Theorem 3 *Consider a (countable) sequence of games $\{\Gamma(\mathscr{I}_{n^*}(T), F^*_{n^*})\}$ for $n^* = 1, 2, ...$ such that each $\mathscr{I}_{n^*}(T)$ is coarser than $\mathscr{I}_{n^*+1}(T)$. Then, as contestants become fully ignorant*
(i) the (unique) optimal success function converges uniformly towards the all-pay auction success function;
(ii) the size of the prize decreases monotonically.

Proof. See Appendix C ■

Theorem 3 is a corollary of Theorem 2: the statements in Theorem 3 are obtained by studying the behavior of the linear piecewise difference-form SF derived in Theorem 2 as information gets coarser. Specifically, notice that for general information settings a contestant's probability of winning the prize increases at the rate $\frac{1}{2\bar{r}k^*(\mathscr{I}_n)}$, until the probability reaches one. Theorem 3 says that the optimal success function converges towards the celebrated all-pay auction SF as the information between the contestants becomes negligeable i.e. they are ignorant of each other. From this, we conclude (ii) i.e. when agents have more information (about each other), the prize to spur competition increases. Intuitively, when each agent has a substantial lack of information on his rival's stream of outputs, the cardinality of the minimal set of discontinuity points becomes smaller. More specifically, inspection of the formulae of $N^*(\mathscr{I})$, – recall that this scalar gives the largest number of periods for which both participants have always some incentives to exert maximal effort in the game under the all-pay auction type –, reveals that refining the information between contestants 'output histories raises the number of discontinuity points of the optimal success function, thereby decreasing its jumps. From this, it follows that, for a decreasing set of discontinuity points this raises the marginal incentives at any history, thereby reducing the prize set by the administrator.

5. Robustness

The purpose of this section is to test the robustness of our results in a "dual" version of the above model where we *fix the prize* and examine the *variable behavior* of agents that is induced by the prize. The multi-period contest model of Dubey and Haimanko (2003) offers an avenue to do this. The game is played as in the above model of contest, expect that each agent is *fully informed of the entire past history of efforts undertaken and outputs realized*. This leads us to view the contest as a tree of *perfect information* modulo simultaneous moves of agents. Let $\Psi(T, F, B)$ be the resulting extensive form game with perfect information (modulo simultaneous moves of agents in each period). We analyze this game for its subgame perfect equilibria (SPE) in behavioral strategies. Following Dubey and Haimanko (2003), we assume the administrator values the expected total output produced by the agents. But it is natural for him to also care about the probability with which he promotes the most skilled agent (skill will be defined in terms of the probability of promotion), provided this does not have an adverse impact on the output. Formally, the organizer's objective is to maximize $U(Q/(2T), P)$, where Q is the expected total output of the agents across the T periods, and P is the probability that the more skilled agent is promoted. We assume that U is strictly increasing in Q and nondecreasing in P. To capture the fact that, beyond a point, output outweighs accurate promotion from the planner's perspective, we assume that there exists $\tau \in (0, 1)$ such that $U(Q, 0) > U(Q - \tau, 1)$ for all $Q \in [\tau, 1]$. This assumption means that the planner's utility function depends on Q and P and that *no improvement in the probability of choosing the more skilled agent can compensate the organizer for the loss of the critical fraction τ of expected total output.*[17] We also suppose the following specifications.

Agents simultaneously choose to either work or shirk (w or s hereafter). As in Dubey and Haimanko, we assume a binary set of outputs $Q = \{0, 1\}$ and conditional on his effort e_i, the probability of getting output q_i is $p(q_i \mid e_i)$, such that $p(q_i = 1 \mid e_i) = p_i$ if $e_i = w$ (resp. 0 if $e_i = s$). Agent 1 is more skilled than agent 2 in the sense that $p_1 > 2p_2$. Let c denote the disutility incurred by

[17] Note that an arbitrarily small "toleration" level τ, reflects the lexicographic character of the planner's preferences over P and Q.

any agent i when he works in all T periods. We suppose that this disutility is uniform and additive across time. Agents' disutilities to work is assumed not too hight relative to their productivity and the value they place on promotion in the sense that $p_2(1-p_1)B > c$. Agent i's disutility $D_i^{h^{T+1}}$ at terminal history, h^{T+1}, is the sum of the one period disutilities. Thus the expected payoff to agent i at the terminal history h^{T+1} is: $P_i^{h^{T+1}} \cdot B - D_i^{h^{T+1}}$ where $P_i^{h^{T+1}}$ is i's probability of promotion at terminal history h^{T+1}. Let $\mathscr{F}_*^{NE}$ be the set of NE-feasible SFs defined for the full information case as in Theorem 1 i.e. the set of NE-feasible SFs when agents have perfect information regarding each other's outputs. If the administrator adopts a SF in $\mathscr{F}_*^{NE}$, we obtain the following result.

Proposition 1 *Suppose $F \in \mathscr{F}_*^{NE}$. Both agents work in every period with probability 1 in any SPE of $\Psi(T, F, B)$, and thus their expected total output is $(p_1 + p_2)T$.*

Proof. See Appendix C ■

A precise characterization of the optimal SF(s) when the administrator is interested in simultaneously maximizing his expected output *while* rewarding almost always the most skilled agent is out of the scope of this paper. In particular, whether a Tullock's functional form might outperform the class of difference-form SFs – when it approaches the all-pay auction – remains an open question.[18] However, consider the following family of (discontinuous) piecewise linear difference-form SFs a la Che and Gale: specifically, we define the set of of the form,

$$F_\epsilon(r) = \begin{cases} 1 - (T-k)\epsilon/2 \text{ if } \ r \in I_{k+1}; \\ \dots; \\ \frac{1}{2} \text{ if } \ r \in I_0; \\ \dots; \\ (T-k)\epsilon/2 \text{ if } \ r \in I_{-(k+1)}, \end{cases}$$

with $k = -\frac{(k^*(\mathscr{I})+1)}{2}, ..., 0, ..., \frac{k^*(\mathscr{I})+1}{2}$ if $k^*(\mathscr{I})$ odd or $k = -(\frac{k^*(\mathscr{I})}{2}), ..., 0, ..., \frac{k^*(\mathscr{I})}{2}$ if $k^*(\mathscr{I})$ even and $\epsilon > 0$. By Proposition 1, under this set of ϵ-all-pay auction SFs, $\mathscr{F}_\epsilon \subset \mathscr{F}_*^{NE}$, all agents work in every period with probability 1 in any SPE of $\Psi(T, F, B)$ (thereby maximizing the utility of an administrator first and foremost concerned about the expected total output that his agents

[18]In this section Theorem 2 does not apply since the value of the prize is assumed to be fixed.

will produce), while for a sufficiently small (positive) ϵ, the prize goes (almost always) to the agent who obtained the *largest output*. This observation immediately reveals that when the planner has also an auxiliary interest in promoting the more skilled agent, the optimal difference-form SF(s) must lie in $\mathscr{F}_\epsilon \subset \mathscr{F}_*^{NE}$, provided T is large enough. This insight follows from Lemma 6.3 of Dubey and Haimanko (2003) which guarantees that the probability that the administrator does not become aware of a lead of $R_1(\mathbf{q}_1(T), \mathbf{q}_2(T)) = \epsilon T$ for agent 1 at a terminal history is *exponentially low* for $0 < \epsilon \leq 1$ (under the all-pay auction SF). Alternatively, the administrator may want to maximize his expected output and detect *the most productive agent* in the contest with the smallest value of promotion i.e. the one who obtains the largest output. In this case, all SFs lying in $\mathscr{F}_\epsilon \subset \mathscr{F}_*^{NE}$, are superior in the class of difference-form SFs. This is so since (by Proposition 1) agents exert maximal effort levels under any SF in $\mathscr{F}_*^{NE}$, and so the expected total output cannot possibly be surpassed under any other SFs.

6. Concluding Remarks

This paper has made a first step towards analyzing the optimal mode of selection of the winner in a dynamic contest where the administrator's major objective is to maximize the expected total output. Our results show that, unlike what is commonly assumed, the optimal scheme is generally a *noisy* ranking system of the *difference-form*: the win probabilities must depend on the difference between the two (stochastic) aggregate outputs of the contestants. This conclusion can even be sharpened if the organizer's objective is to make contestants work at their maximal effort levels. In this case, the optimal ranking scheme corresponds to (a generalized version of) the piecewise linear difference-form success functions introduced by Che and Gale (2000). Characterizing optimal success functions for more general class of objective functions is a promising avenue for further inquiry.

Appendix A

We first present the main notations we use to prove our results.

Let $\mathcal{I}_i(t^{'}) \subset \mathcal{I}_i$ be the set of all possible information sets for agent i from period $t = t^{'}, ..., T$. Throughout the Appendix, we denote by $\sigma_{i|t^{'}}$ the strategy that σ_i induces at the set of information sets, $\mathcal{I}_i(t^{'})$. For ease of exposition, throughout the Appendix, we only use notations for the case where the Q_i's are two infinite sets of performance.

It is convenient to define,

$$R_i(t) := \Big\{\delta \in \mathbb{R} : \text{ it exists } ((q_i(\tau))_{\tau=t+1}^T, (q_{-i}(\tau))_{\tau=t}^T) \in Q_i^{T-t-1} \times Q_{-i}^{T-t} : \sum_{\tau=t+1}^{T} q_i(\tau) - \sum_{\tau=t}^{T} q_{-i}(\tau) = \delta\Big\}.$$

Hence, the set $R_i(t)$ is nothing but the set of leads consistent with vectors of output starting at $t+1$ for agent i and at t for agent $-i$,

$$(q_i(t+1), ..., q_i(T), q_{-i}(t), q_{-i}(t+2), ..., q_{-i}(T)) \in Q_i^{T-(t+1)} \times Q_{-i}^{T-t},$$

leaving the performance at t of agent i *unspecified* and resulting in a difference *for agent i* of δ. Let $x_i - y_{-i} \in \left[-(\overline{q}_{-i} - \underline{q}_i), \overline{q}_{-i} - \underline{q}_i\right]$ be any possible lead that may be attained by i within one period. Given an output history $(\mathbf{q}_1(t-1), \mathbf{q}_2(t-1))$ at t, we define

$$\Delta(\mathscr{I}_{n_i}, t) = \begin{cases} q_i(t) - q_{-i}(t) \text{ if } \ i \text{ observes } -i\text{'s output at } t; \\ x_i - y_{-i} \text{ otherwise,} \end{cases}$$

as a **possible lead** for i at period t which is consistent with his information partition $\mathscr{I}_{n_i}$. Accordingly, $r_i(h^t)$ is a (the unique) possible lead[19] of i at history h^t consistent with his information set, $\mathcal{I}_i(h^t) \in \mathscr{I}_{n_i}$. Thus, $r_i(h^t) \equiv \sum_{\tau \notin N_i(t)} \Delta(\mathscr{I}_{n_i}, \tau) + \sum_{\tau \in N_i(t)} \Delta(\mathscr{I}_{n_i}, \tau)$ denotes a (the unique) possible lead of i at $\mathcal{I}_i(h^t)$ when i perfectly observes his rival's output at periods $\tau \in N_i(t)$.[20] Conditional on $r_i(h^t)$, the profile σ and the *value of the random variable* $X_{T-t}^i(\sigma) = \delta$ (i.e. the value taken on by the lead over the next $T-t$ periods but without the performance of i at period t evaluated from period t to T that leaves i's output *unspecified* at t), the **δ-conditional expected utility** to win the prize for contestant i at his information set $\mathcal{I}_i(h^t)$ is

$$\int_{Q_i} u_F \circ B(r_i(h^t) + q + \delta) dG_i(q \mid e).$$

Denote by $G(\cdot \mid \sigma_{i|t}, \sigma_{-i|t+1})$, the joint distribution function over the leads conditional on the profile $(\sigma_{i|t+1}^*, \sigma_{-i|t}^*)$. Then, conditional on $r_i(h^t)$, the profile σ, the **overall expected utility** to win the prize for agent i at his information set $\mathcal{I}_i(h^t)$ is

$$\int_{R_i(t)} \left[\int_{Q_i} u_F \circ B(r(h^t) + q + \delta) G_i(q \mid e)\right] dG(\delta \mid \sigma_{i|t}^*, \sigma_{-i|t+1}^*).$$

[19] This arises when i observes everything except his rival's efforts.

[20] Recall that $N_i(t)$ denotes the set of all periods $1 \le \tau \le t$ where i can *perfectly observes his rival's output.*

where $u(B(r)F(r) \equiv u_F \circ B(r)$. Conditional on $r_i(h^t)$, the profile σ and the *value of the random variable* $X^i_{T-t}(\sigma) = \delta$, the **incentive** of i to exert maximal effort level, $\sigma_i(\mathcal{I}_i(h^t)) = e^*$ rather than $e \neq e^*$ at information set $\mathcal{I}_i(h^t) \in \mathscr{I}_{n_i}$ is

$$\int_{R_i(t)} \left[\int_{Q_i} u_F \circ B(r_i(h^t) + q + \delta) d\Delta G_i(q \mid e) \right] dG(\delta \mid \sigma^*_{i|t}, \sigma^*_{-i|t+1})$$

where $\Delta G_i(q \mid e) \equiv G_i(q \mid e^*) - G_i(q \mid e)$ for all $q \in Q_i$ and $e \neq e^*$. Thereafter we denote the above equation by $\Delta((F, B) \,| r_i(h^t))$.

For ease of exposition, the proof of the properties of NE-feasible policies ($u_{i\,|F_R} \circ B_R : \mathcal{Q}_R \to \mathbb{R}_{++}$ is (non constant) increasing over each $I_r \in \mathscr{P}^*(\mathscr{I})$), is given for the case of difference-form policies ,(F_R, B_R), when the rule R corresponds to the *special additive rule*, $R_i(\mathbf{q}_i(T+1), \mathbf{q}_{-i}(T+1)) := \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$. We denote such policies defined on this rule by (F, B). But all our results hold for any other *additive* rule with obvious modifications. Then, we complete the proof of Theorem 1 and shows that *only difference-form policies* are indeed NE-feasible.

Proof of Theorem 1

($\Rightarrow$). Since we want to implement σ^* as a NE, we must necessarily have that $\Delta((F, B) \,| r_i(h^t)) > 0$ for all $\mathcal{I}_i(h^t) \in \mathscr{I}_{n_i}$. We first show that a necessary and sufficient condition for this property is that (F, B) be (non-constant) increasing on each $\{r_i(h^t) + \delta + q_i : (\delta, q_i) \in A_i(t) \times Q_i\} \equiv Q_i(r_i(h^t))$. Let $\int_{Q_i} u_F \circ B(r_i(h^t) + q + \delta) d\Delta G_i(q \mid e) \equiv \Delta(F \,| r_i(h^t) + \delta)$ and denote the restriction of $u_F \circ B$ to $Q_i(r_i(h^t))$ by $u_F \circ B_{|Q_i(r_i(h^t))}$. By a well-known property of first-order stochastic dominance, we have that $G_i(\cdot \mid e^*)$ strictly first-order dominates $G_i(\cdot \mid e)$ for $e \neq e^*$ if and only if, for any non-constant increasing function $u_F \circ B_{|Q_i(r_i(h^t))} : Q_i(r_i(h^t)) \to \mathbb{R}_{++}$ the inequality $\Delta(F \,| r_i(h^t) + \delta) > 0$ holds. Hence, when this property is met, we have that $\Delta((F, B) \,| r_i(h^t) + \delta)) > 0$ holds for any $\delta \in R_i(t)$. It is then readily shown that $\Delta((F, B) \,| I_i(r(h^t))) > 0$ for any h^t if and only if $B(F) < \infty$. This follows easily from the use of the assumptions made on the utility functions i.e. $u(B) \to \infty$ as $B \to \infty$ and the disutility of effort i.e. $D_i(\mathbf{e}^*) - D_i(\mathbf{e})$ is bounded above for all $\mathbf{e} \in E^T \setminus \{\mathbf{e}^*\}$.

($\Leftarrow$). By contraposition. Suppose $u_F \circ B$ is non-constant decreasing on some $Q_i(r_i(h^T))$ while (F, B) is a NE-feasible policy. Then, this implies that conditional on *any* $q_{-i} \in Q_{-i}$, i has no incentives to work at the last period since $\int_{Q_i} u_F \circ B(r_i(h^T) + q_i - q_{-i}) d\Delta G_i(q_i \mid e) \leq 0$. By definition this implies that $\Delta((F, B) \,\big| r_i(h^T)) \leq 0$. Since $D_i(\mathbf{e}^*) - D_i(\mathbf{e}) > 0$, for all $\mathbf{e} \in E^T \setminus \{\mathbf{e}^*\}$, we have an obvious contradiction. ■

Completion of Theorem 1: Only difference-form policies are NE-feasible

We proceed by contradiction. Suppose that R_i does *not* represent the order $\succeq_i^{\mathscr{I}}$. This implies that there exists two terminal histories, $(\mathbf{q}_i(T), \mathbf{q}_{-i}(T))$ and $(\mathbf{q}_i^{'}(T+1), \mathbf{q}_{-i}^{'}(T))$ that could be reached from $r_i(h^T)$ with $(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) \succ_i^{\mathscr{I}} (\mathbf{q}_i^{'}(T+1), \mathbf{q}_{-i}^{'}(T+1))$ while agent i's probability of winning at these terminal history are such that $F \circ R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) \leq F \circ R_i(\mathbf{q}_i^{'}(T), \mathbf{q}_{-i}^{'}(T))$. Hence, there exists a last period history, h^T with $r_i(h^T) \in [-k^*(\mathscr{I})\overline{r}, k^*(\mathscr{I})\overline{r}]$ such that, $F \circ R_i \equiv \psi_i$ is weakly decreasing on the interval $[r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]$. This violates the conditions of Theorem 1 since the function $u_{i|F\circ R_i} \circ \overline{B} : [-T\overline{r}, T\overline{r}] \rightarrow \mathbb{R}_{++}$ is no longer (non constant) increasing over each $I_r \in \mathscr{P}^*(\mathscr{I})$. It follows that if R_i does not represents the order $\succeq_i^{\mathscr{I}}$ on $Q_1^T \times Q_2^T$, then σ^* cannot be a NE of $\Gamma(\mathscr{I}, F_R, \overline{B}_R)$ for any finite bonus $\overline{B}$. An obvious contradiction.■

Appendix B (Optimal policies)

Recall that we simply denote policies (F_R, B_R) by (F, B) when the rule R corresponds to the *special additive rule*, $R_i(\mathbf{q}_i(T), \mathbf{q}_{-i}(T)) := \sum_{\tau=1}^{T} q_i(\tau) - q_{-i}(\tau)$. In what follows, we show the optimality of picking a reward policy $(F, \overline{B})$ such that $\overline{B}$ is the constant function over all other policies using a non constant prize function. To do so, we first prove Lemma 2 which shows that conditional on a lead $r_i(h^T)$ reached at the last period, the agents' incentives to exert e^* at $I_i(h^T)$ is increasing in the jumps assigned on$[-r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]$. From this, we immediately conclude that it is always optimal to *minimize* the number of jumps. Moreover, since σ^* induces *intertemporal equality of marginal disutility* for all agents, we deduce that all the conditional marginal increments in the probability of winning (by exerting e^* rather than $e \neq e^*$) are maximized if and only the rate of increase of the SF is *uniform* over each interval, $[-r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]$. From this we will deduce all the ensuing properties that an optimal success function must meet.

Local property of the SF. First, we show that *locally*, the optimal SF will be a step function. Then, iterating this reasoning, we infer some global properties about F^* on its entire domain. Consider the form of the optimal policy (F^*, B^*) when B^* is a constant function. First, note that minimizing the prize implies the maximization of each agent's incentive to exert $\sigma_i^*(\mathcal{I}_i(h^t)) = e^*$ for all $\mathcal{I}_i(h^t) \in \mathscr{I}_{n_i}$. However, it is easy to see that the worst information sets to get contestants to work, is at the last period (if they work here, they will work at all the other information sets). Hence we must consider the maximization of incentives at information sets, $\mathcal{I}_i(h^T)$ i;e., $\max \Delta(F \left| r_i(h^T) \right.)$. W.lo.g., consider the restriction of F^* to $Q_i(r_i(h^T)) \subset [0, k^*(\mathscr{I})\overline{r}]$ which we denote by $F^*_{|Q_i(r_i(h^T))}$. For each such restriction, $\Delta(F^*_{|Q_i(r_i(h^T)} \left| r_i(h^T) \right.)$, the maximand is linear in the choice variable so the maximizer is a step function, $F^*_{z^* \,|Q_i(r_i(h^T))}$ of the form

$$F_{z^* \,|Q_i(r_i(h^T))}(r) = \begin{cases} 1 \text{ if } & r \geq z^*; \\ \frac{1}{2} \text{ if } & r < z^* \end{cases}$$

with $z^* = \arg\max_{\widehat{z} \in [-r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]} \Delta(F^*_{\widehat{z}} \left| r_i(h^T) \right.)$. Since F^* is symmetric around the origin, it is clear that $\mu_{F^*}(z^*(r_i(h^T) + q)) = \mu_{F^*}(-z^*(r_i(h^T) + q))$. Using the full support assumption of the cdf $G(\cdot \mid e^*)$ and the assumption of infinite set of outputs, we have that for a given $\mathscr{I}$, *any* lead in $[-\overline{r}(k^*(\mathscr{I})), \overline{r}(k^*(\mathscr{I}))]$ can be reached with a strict positive probability at the last period. This means that $F_{z^* \,|Q_i(r_i(h^T))}$ must ideally be true for each interval centered at $r_i(h^T) \in [-\overline{r}(k^*(\mathscr{I})), \overline{r}(k^*(\mathscr{I}))]$ of length $2\overline{r}$. (this follows from the assumption made on the contestants' sets of outputs). Considering the constraint of symmetry ($F^* \in \mathscr{F}$) and the fact that this must be true for each $r_i(h^T)$, this entails that the jump at each discontinuity point of F^* must be set as large as possible since one wants to minimize the prize and the number of jumps is optimally minimized (under the constraint imposed by Theorem 1). By extending this result on the entire domain of F, we conclude from the earlier observations that under NE-feasible SFs, all the conditional marginal increments in the probability of winning must optimally be constant positive on *any* interval $[r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}] \subset [-k^*(\mathscr{I})\overline{r}, k^*(\mathscr{I})\overline{r}]$. ■

The next result shows that we can restrict attention to reward policies of the type $(F, \overline{B})$ with $F \in \mathscr{F}^{NE}$.

Proof of Lemma 1 Suppose B is a non constant function. We know that conditional on any $r_i(h^T)$ the maximizers of i's incentive, (F^*, B^*), will be two step functions. Hence, under such a policy, $B^*(F^*) \to \infty$ so as to maximize the jump at each $r_i(h^T)$. Clearly, for the administrator, there always exists another policy, $(F^{'}, \overline{B})$ with $F^{'} \in \mathscr{F}^{NE}$ and $\overline{B}$ a constant prize function such that $\overline{B} \equiv B(F^{'}) < \infty$ which proves that (F^*, B^*) cannot be an optimal policy.■

Using Part 1 and Lemma 1, the next result provide a characterization of the set of potential optimal policies.

Proof of theorem 2 We show that (given a partition $\mathscr{I}$) (and when cdfs satisfy first order stochastic dominance), then there is a unique optimal SF which is the *continuous* piecewise linear increasing SF F^*

$$F^*(r) = \max\left\{\min\left\{\frac{1}{2} + \frac{r}{2k^*(\mathscr{I})\overline{r}}, 1\right\}, 0\right\}.$$

By the full support assumption of the cdfs $G_i(\cdot \mid e)$ and the fact each Q_i is a closed interval of $\mathbb{R}_+$, observe that all leads in $[-(k^*(\mathscr{I})-1)\overline{r}, (k^*(\mathscr{I})-1)\overline{r}]$ can be reached with a strict positive probability at the last period. Pick two sets of discontinuity points $\{z_i\}$ (resp. $\{z_i^{'}\}$) taking their values in interval $[r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]$ (resp. $[r_i(h^{'T}) - \overline{r}, r_i(h^{'T}) + \overline{r}]$). The optimal SF F^* must possess the following property:

(P1) conditional on any $r_i(h^T)$, all intervals $[r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r}]$ (resp. $r_i(h^{'T}) \in [r_i(h^{'T}) - \overline{r}, r_i(h^{'T}) + \overline{r}]$) must have been exactly assigned the *same mass* i.e. $\sum_{i=1}^{n} \mu_{F_n^*}(z_i) = \sum_{i=1}^{n} \mu_{F_n^*}(z_i^{'})$.

Property $(P1)$ is a direct consequence of Lemma 2.

Proof of Lemma 2

We proceed by contradiction. If F^* does *not* increase at a constant rate, there exist two histories h^T and $h^{'T}$ such that $F^*(r_i(h^T)+\overline{r}+\epsilon) - F^*(r_i(h^T)+\overline{r}) \neq F^*(r_i(h^{'T}) - \overline{r} + \epsilon) - F^*(r_i(h^{'T}) - \overline{r})$. But then, this would imply that i's last period incentives at those histories are such that $\Delta(F_n^* \left| r_i(h^T)\right.) \neq \Delta(F_n^* \left| r_i(h^{'T})\right.)$. Thus, one could decrease some jumps at some histories to increase the jumps at other histories, thereby reducing the size of the prize at which all agents exert unconditional maximal effort level. Therefore, this last statement would contradict the optimality of F^*. This shows that F^* will necessarily increase at a *constant rate*.■

Continuity of the optimal SF.

We now show that $F^* \in \mathscr{F}^{NE}$ is necessarily continuous. To see this, it suffices to note that in this case $\Delta(F^* \left| r_i(h^T)\right.) = \Delta(F^* \left| r_i(h^{'T})\right.)$ for all $r_i(h^T), r_i(h^{'T})$ since the rate of increase is uniform on the entire domain of F^* (this directly follows from the definition of a linear continuous function).

Next, we show that if $\Delta(F^* \left| r_i(h^T)\right.) = \Delta(F^* \left| r_i(h^{'T})\right.)$ for all $r_i(h^T), r_i(h^{'T})$, then F^* is continuous. We proceed by contradiction. Let,

$$z_i^*(r_i(h^T) = \arg \max_{\widehat{z} \in [-r_i(h^T) - \overline{r}, r_i(h^T) + \overline{r})} \Delta(F_{\widehat{z}}^* \left| r_i(h^T)\right.),$$

and $z_i^*(r_i(h^{'T})$ defined similarly. Suppose $F^* \in \mathscr{F}^{NE}$ satisfies P1 above but is not continuous. Hence, F^* has some jumps of discontinuity and there exists by definition $z_i^*(r_i(h^T) \neq z_i^*(r_i(h^{'T})$. From this, we conclude that $\Delta(F^* \left| r_i(h^T)\right.) \neq \Delta(F^* \left| r_i(h^{'T})\right.)$ if F^* has all its jumps uniform (which must be the case by P1). Hence, one can increase the jumps at leads,

$$r_i^*(h^T) = \arg \min_{\widehat{r}_i(h^T)\in[k^*(\mathscr{I})\overline{r},k^*(\mathscr{I})\overline{r}]} \Delta(F^*_{\widehat{z}} \left| \widehat{r}_i(h^T)\right.),$$

by reducing the mass at leads $r_i^*(h^T)$ maximizing the one period incentive, thereby decreasing the prize at which agents exert σ^*. This implies that we have found another SF, $F \neq F^*$, with $B(F) < B(F^*)$ contradicting the optimality of F^*. This completes the proof.■

Appendix C

Proof of Theorem 3: Coarser information partitions

We prove Theorem 3 for any NE-feasible SFs that is linear piecewise difference-form but *not* necessarily continuous. We first state the following Lemma.

Lemma Given $\mathscr{I}_{n^*}$, the largest number of periods wherein contestants shirk, $k(\mathscr{I}_{n^*})$, is decreasing in $n^* = \min\{n_1, n_2\}$.

Proof. We shall use the notations introduced in Section 3. We show that $N^*(\mathscr{I})$ is increasing in n^*. First, note that for a given k, the optimistic scenario, $\overline{h}(k, \mathscr{I}_{n^*})$ is increasing in n^*. This follows since the number of observable periods t where $\Delta^O(\mathscr{I}_{n_i}, t) < 0$ decreases. It is also easy to see that $\overline{N}(\mathscr{I}_{n^*})$ is increasing in n^*. A symmetric argument shows that $\overline{N}(\mathscr{I}_{n^*})$ is also increasing in n^*. Thus, we have that $N^*(\mathscr{I}_{n^*})$ is increasing in n^*. Since this is true for both contestants, $N^*(\mathscr{I}_{n^*})$ is also increasing. Hence, $T - N^*(\mathscr{I}_{n^*}) \equiv k(\mathscr{I}_{n^*})$ decreases as n^* increases, which completes the proof.■

With the Lemma above, we complete Theorem 3 as follows. *(i)* Define

$$F^*_{n^*}(r) = \max\left\{\min\left\{\frac{1}{2} + \frac{k}{T-n^*-1}, 1\right\}, 0\right\} \text{ if } r \in I_k.$$

Let $\|F_{n^*}(\cdot) - F(\cdot)\| \equiv \sup_{\widehat{k}} \left|F_{n^*}(\widehat{k}) - F(\widehat{k})\right|$. In general, $F_{n^*}(\cdot)$ is a step function $F_{n^*}(\cdot) - F(\cdot)$. Hence, it attains its supremum at a point where $F_{n^*}(\cdot)$ jumps. It is then easy to observe that $F^*_{n^*}(\cdot)$ converges uniformly to the all-pay auction.

(ii) Pick a sequence of games, $\{\Gamma(\mathscr{I}_k, F_k^*)\}$ for $k = 1, 2, ...$ such that contestants becomes ignorant. Using formulae (1), it is immediate that $k^*(\mathscr{I})$ converges monotonically to 0. Hence, the (uniform) mass assigned to each

jump increases monotonically, thereby increasing at any information set the contestants' incentives. Since this is true at any point of discontinuity of F^*, this implies that incentive increases at all information sets. As an immediate consequence, this means that σ^* can be sustained as a NE with a smaller prize.■

Proof of Proposition 1

We will show by backward induction that in every period both agents work unconditionally in any SPE of $\Gamma(F, B)$. Let h^T be an history in the last period T and let i be an agent whose total output along the path leading from h^0 to h^T has a lead of $r_i(h^T)$ units. We claim that the change in i's payoff, when he switches from shirk to work at h^T, is positive, no matter what action the other agent j chooses at h^T. Given $r_i(h^T)$, **if j shirks** at h^T then, by working at h^T, i ensures the payoff of $p_i F(r_i(h^T)+1)B + (1-p_i)F(r_i(h^T))B$, compared to $BF(r_i(h^T))$ when he shirks. Thus, if F is a strictly increasing function, $B(p_i(F(r_i(h^T)+1) - F(r_i(h^T)))) - c$ is the positive gain in payoff for i in this case (now taking the costs of work also into account).

If j works at h^T then, by working at h^T, i ensures the payoff of, $p_i(1-p_j)F(r_i(h^T)+1)B + (1-p_i)(1-p_j)F(r_i(h^T))B + BF(r_i(h^T)-1)(1-p_i)p_j + BF(r_i(h^T))p_ip_j$, compared to $(1-p_j)F(r_i(h^T))B + p_jF(r_i(h^T)-1)B$ when he shirks. Conditional on $r_i(h^T)$, the switch from work to shirk leads to a gain in payoff of (now taking the costs of work also into account),

$$B\left[p_i(1-p_j)(F(r_i(h^T)+1) - F(r_i(h^T))) + (F(r_i(h^T) - F(r_i(h^T)-1)))p_ip_j\right] - c,$$

which is positive when F is a strictly increasing function. Thus the change in the **unconditional payoff** (now taking the costs of work also into account) if agent i switches from shirk to work at **any** history h^T is always positive **whatever is the action of** j at h^T may be. Therefore i's SPE strategy will tell him to work at h^T with probability 1. Given that decision by i, the other agent j will also work at h^T with probability 1. Thus the SPE strategy of j tells him to work unconditionally at h^T. Now, assume that h^t is an history in period $t < T$, and that the agents work at every history following h^t. No matter what agents' actions at h^t may be, on subsequent periods both agents work by the inductive assumption. Choose any $r_i(h^t)$ such that there exists a path of outputs leading to $r_i(h^T)$. Conditional on $r_i(h^t)$, as in the above paragraph, the change in payoff when agent i switches from shirk to work at h^t is again always positive. Therefore the SPE strategies of both agents tell them to unconditionally work at h^t. This finishes the inductive proof.■

References

Alcade, J. and M. Dahm (2007) "Tullock and Hirshleifer: a meeting of the minds" *Review of Economic Design* **11**, 101-124.

Baik, K.H. (1998) "Difference-form contest success functions and effort level in contests" *European Journal of Political Economy* **14**, 685-701.

Che, Y.K. and I. Gale (1998) "Caps on political lobbying" *American Economic Review* **88**, 643-651.

Che, Y.K. and I. Gale (2000) "Difference-form contests and the robustness of all-pay auctions" *Games and Economic Behavior* **30**, 22-43.

Corchón, L. (2007) "The theory of contests: a survey" *Review of Economic Design* **11**, 69-100.

Corchón L. and M. Dahm (2010) "Foundations for contest success functions" *Economic Theory* **43**, 81-98.

Corchón, L. and M. Dahm (2011) "Welfare maximizing contest success functions when the planner cannot commit" *Journal of Mathematical Economics* **47**, 309-317.

Dasgupta, A. and K.O. Nti (1998) "Designing optimal contests" *European Journal of Political Economy* **14**, 587-603.

Demski, J. S., and D. Sappington (1984) "Optimal incentive schemes with multiple agents" *Journal of Economic Theory* **33**, 152-171.

Dubey, P. and O. Haimanko (2003) "Optimal scrutiny in multi-period promotion contests" *Games and Economic Behavior* **42**, 1-24.

Dubey, P. and C.W. Wu (2001) "Competitive prizes: when less scrutiny induces more effort" *Journal of Mathematical Economics* **36**, 311-336.

Epstein, G.S. and S. Nitzan (2006) "The politics of randomness" *Social Choice and Welfare* **27**, 423-433.

Fu, Q. and J. Lu (2009) "The optimal multi-stage contest" *Economic Theory* (forthcoming).

Fu, Q. and J. Lu (2012) "Micro foundations of multi-prize lottery contests:

a perspective of noisy performance ranking" *Social Choice and Welfare* **38**, 497-517.

Gradstein, M. and K.A. Konrad (1999) "Orchestrating rent seeking contests" *Economic Journal* **109**, 535-545.

Green, J. and N. Stokey (1983) "A comparison of tournaments and contracts" *Journal of Political Economy* **91**, 349-364.

Hillman, A.L. and J.G. Riley (1989) "Politically contestable rents and transfers" *Economics & Politics* **1**, 17-40.

Hirshleifer, J. (1989) "Conflict and rent-seeking success functions: ratio vs. difference models of relative success" *Public Choice* **63**, 101-112.

Itoh, H. (1991) "Incentives to help in multi-agent situations" *Econometrica* **59**, 611-637.

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

Konrad, K.A., Kovenock, D (2009) "Multi-battle contests" *Games and Economic Behavior* **66**, 256-274.

Konrad, K.A., Kovenock, D (2010) "Contests with stochastic abilities" *Economic Inquiry* **48**, 89-103.

Lazear, E.P. and S. Rosen (1981) "Rank order contest as optimum labor contracts" *Journal of Political Economy* **89**, 841-864.

Lockwood, B (2000) "Production externalities and two-way distortion in principal-multi-agent problems" *Journal of Economic Theory* **92**, 142-166.

Münster, J (2009) "Repeated Contests with Asymmetric Information" *Journal of Public Economic Theory* **11**, 89-118.

Mookherjee, D (1984) "Optimal incentive schemes with many agents" *Review of Economic Studies* **51**, 433-446.

Nalebuf, B Stiglitz, J. (1983) "Prizes and incentives: Towards a general theory of compensation and competition" *Bell Journal of Economics*, **14**, 21-43.

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 3-15