

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1254

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

OPTIMAL SCRUTINY IN MULTI-PERIOD PROMOTION TOURNAMENTS

Pradeep Dubey and Ori Haimanko

April 2000

Optimal Scrutiny in Multi-Period Promotion Tournaments

Pradeep Dubey* and Ori Haimanko†

April 2000

Abstract

Consider a principal who hires heterogeneous agents to work for him over T periods, without prior knowledge of their respective skills, and intends to promote one of them at the end. In each period the agents choose effort levels and produce random outputs, independently of each other, and are fully informed of the past history of outputs.

The principal's major objective is to maximize the total expected output, but he may also put some weight on detecting the higher-skilled agent for promotion. To this end, he randomly samples n out of the T periods and awards the promotion to the agent who produces more on the sample. This determines an extensive form game $\Gamma(T, n)$, which we analyze for its subgame perfect equilibria in behavioral strategies.

We show that the principal will do best to always choose a *small* sample size n . More precisely, if $\eta(T)$ is the maximal optimal sample size, then $\eta(T)/T \rightarrow 0$ as $T \rightarrow \infty$.

Journal of Economic Literature Classification Numbers: C72, D82, J41.

Key words: multi-period promotion tournaments, extensive form games, subgame perfect equilibria, undominated sample sizes.

*Center for Game Theory in Economics, SUNY at Stony Brook, NY 11794. E-mail: pradeepkdubey@yahoo.com

†Cowles Foundation for Research in Economics, Yale University, P.O.Box 20-8281, New Haven, CT 06520-8281. E-mail: ori.haimanko@yale.edu

1 Introduction

The literature on tournaments considers a principal who awards prizes to his agents according to the relative ranking of their outputs (see, e.g., Lazear and Rosen (1981), Green and Stokey (1983), Nalebuff and Stiglitz (1983), Rosen (1986)). The objective of the principal is to maximize total expected output net of the prizes he hands out. Agents, on the other hand, have a natural disutility for effort and the prizes must be sufficiently lucrative in order to induce them to work. The output of each agent is stochastic but positively correlated with his effort level.

While the prize awarded to an agent is highly dependent on his rivals' outputs, his effort level itself is mostly assumed to be unconditional on what they may be doing. In effect, tournaments are viewed as one-shot strategic games in which all agents choose effort levels simultaneously and independently of each other. However, there are many situations in which agents in fact compete over a long run prior to being evaluated for awards, and are able to observe each other as the competition unfolds over time. The one-shot model is not appropriate here. Agents' behavior tends to get conditioned quite delicately on the observed history of past outcomes. For instance, an agent would see no reason to work if he had sufficient lead over his rivals and was close to the end of the tournament. But, if he lagged behind, he would carefully estimate his chances of overtaking others before deciding how much effort to put in, and strive hard if there remained sufficient hope to win. Our concern in this paper is with long-run tournaments which are played over T periods by highly observant agents. To make room for sophisticated dynamic behavior, we model such tournaments as T -period extensive form games with perfect information (modulo simultaneous moves of agents in each period).

We consider agents with heterogeneous skills (measured, for any agent, by the maximal output he can produce on average in a period), and focus on T -period tournaments in which the principal awards a prize only to the best-performing agent. Since such tournaments most frequently occur in hierarchical organizations, where agents compete with their peers on a daily basis for a "prize" - the coveted promotion to the next higher echelon, - we call them *promotion tournaments*. We assume that the principal is constrained to make his decision based solely on the observation of agents' outputs. This could be because he is not aware of agents' individual characteristics such as their skills, or simply because the law requires awards to be commensurate with performance and to be based on outputs alone, with no other form of

discrimination permitted.

In any outcome of a T -period promotion tournament each agent produces a stream of outputs over time. A very interesting question, which was hidden from view in the one-shot model, now comes to the surface. How should the principal compare the different streams to decide on the winner? At first glance the aggregate outputs across all T periods seem to provide the natural criterion. But, if there are costs to monitoring, this may be infeasible. In this event one is tempted to think of sampling: the principal could randomly sample n out of the T periods, and promote the agent with the highest output on the sample. We argue here that sampling should indeed be introduced, but for entirely different reasons. Even if the entire stream of outputs is susceptible of costless observation by the principal, he would do best to choose a small sample size n .

The specification of a sample size n determines agents' probabilities of winning the promotion, and thereby their payoffs, for any profile of strategies that they employ in a T -period promotion tournament. This allows us to view the tournament as an extensive form game $\Gamma(T, n)$, which we analyze for its subgame perfect strategic equilibria (PSE) in behavioral strategies. The principal, as was said, values the total expected output produced by agents. But it is natural for him to also care about the probability with which he promotes the strongest (most skilled) agent, provided this does not reduce output to zero. We therefore describe the principal's utility as a function of both variables. It induces a preference relation on the set of PSE of all T -period promotion tournaments. From the principal's perspective, it is evident that sample size k dominates sample size n if each PSE of $\Gamma(T, k)$ is preferred to every PSE of $\Gamma(T, n)$. In Theorem 1 of our paper we prove that if $\eta(T)$ is the largest undominated sample size, then $\frac{\eta(T)}{T} \rightarrow 0$ as $T \rightarrow \infty$, i.e., $\eta(T)$ is a vanishing fraction of T .

In fact, we show somewhat more. Sample sizes which are non-vanishing fractions of T turn out to be quite disastrous asymptotically in T : they lead to expected outputs that are vanishing fractions of T (Theorem 2). The intuition behind this result lies in the fact that the strongest (most skilled) agent can develop a noticeable lead over others in the very beginning of the tournament, by doing relatively little work (on a vanishing fraction of the T periods). Any large sample size almost surely makes the principal aware of this lead. This causes all agents to abandon work at later periods in PSE, as shown in Lemma 9 using backward induction arguments.

In contrast, a better outcome is engendered for the principal through sample sizes 1 or 2. The strongest agent is then forced to work in at least some non-vanishing fraction of the T periods, for otherwise there is a high probability that the principal will fail to take notice of his superior skills (Proposition 3). Moreover, even the weakest agent now stands a reasonable chance to win by working at full capacity throughout the T periods, since his performance on a small sample may far exceed his average. Thus, if the agents value promotion highly enough, they will exert maximal effort in all periods (Propositions 4 and 5). Yet, this does not necessarily imply that sampling once or twice is optimal: the principal may choose a larger sample size, if it does not reduce the total output by much, since the probability of detecting the strongest agent increases significantly. However, the principal must sample on a vanishing fraction of the T periods, otherwise the output falls dramatically by Theorem 2.

Our main results, Theorems 1 and 2, depend crucially on the subgame perfection of the strategic equilibria. Without subgame perfection several other equilibria can be sustained. For instance, if the sample size is large, inactive rivals may force the strong agent to work on a constant fraction of all periods, more than he needs to establish a safe lead over them: for if he did anything else, they would become active and “punish” him by working forever after detecting his deviation one period later, and considerably reduce his probability to win. By the use of more sophisticated threat strategies other equilibria can be exhibited, in which all agents work for a constant fraction of time, provided the sample size is large enough. These equilibria, however, are built with the help of “incredible” threats, which are ruled out by subgame perfection.

The fact that less monitoring by the principal can sometimes elicit more work from agents was pointed out by Cowen and Glazer (1996). Their model is quite different from ours. It has just one agent. The principal offers him a contract specified by means of two parameters: a “threshold” and a “level of scrutiny”¹. The agent must choose a uniform “shirking rate” across all periods. He wins the prize if he is observed to work a fraction of times that exceeds the threshold. For any fixed level of scrutiny, the probability of winning declines as the shirking rate increases, defining an “opportunity

¹More precisely, countably many periods are scrutinized according to a Poisson distribution with mean M , where M is the level of scrutiny (i.e., asymptotically in T , every day is scrutinized with probability $\frac{M}{T}$).

curve” on which the agent optimizes. When the level of scrutiny is raised, the opportunity curve falls for the most part, i.e., this probability falls for most shirking rates. Then it can happen, for *some* indifference curves (between the probability of winning and the shirking rate), that the agent optimizes by shirking more when the level of scrutiny goes up.

The contrast with our model is clear. The need for reduced scrutiny holds in our model in general (for sufficiently large T), not just for suitably chosen configurations of preferences. Moreover, it arises fundamentally from the competition between the agents who choose complicated non-stationary strategies at equilibrium. To interpret their result within our framework, one could imagine a “partial” equilibrium scenario: the agent chooses a stationary “best reply” to an imaginary opponent who is presumed to produce the threshold amount in every period.

In a companion paper, Dubey and Wu (2000), the behavior of agents is fixed at maximal effort levels and variable prizes that induce this behavior are examined. This complements our approach, in which we fix the prize and examine the variable behavior induced by the prize. The analysis in Dubey and Wu (2000) is geared towards games with coarse information and concerns the effects of refining the information, while we deal with perfect information. But in both cases the main theme remains intact: optimal sample sizes of the principal must be small.

The paper is organized as follows. We present our model in Section 2, and state the main results in Section 3. Generalizations, variations, and open problems are in Section 4, and proofs are in the last Section 5.

2 The Model

For ease of exposition, we present here a stripped down version of the model, with just two agents, and defer all generalities to Section 4.

A T -period *promotion tournament* is conducted as follows. At the start of any period $t \in \{1, \dots, T\}$ agents 1 and 2 simultaneously choose to either work or shirk. We assume that if agent i shirks, his output is 0; but, if he works, there is a move of chance which selects output 1 with probability p_i and output 0 with probability $1 - p_i$. Once the outputs of agents are determined, the tournament enters period $t + 1$, when agents simultaneously move again (unless $t = T$, in which case the game ends). At any time, each agent is fully informed of the entire past history of efforts undertaken and outputs

realized. Thus the tournament may be viewed as a tree of perfect information with simultaneous moves. Figure 1 shows the section of the tree when agents choose a pair of effort levels $(e_1, e_2) \in \{work, shirk\}^2$ at the node ω , and then chance selects outputs $(q_1, q_2) \in \{0, 1\}^2$, leading to node $\bar{\omega}$ in the next period:




Figure 1

(The pair (q_1, q_2) is chosen with probability $p_1^{e_1}(q_1) \cdot p_2^{e_2}(q_2)$, where $p_i^{e_i}(q_i) = 1$ if $e_i = shirk$ and $q_i = 0$; and $p_i^{e_i}(q_i) = p_i$ if $e_i = work$ and $q_i = 1$.) This completes the description of the extensive form of the game. It still remains to define agents' payoffs at the terminal nodes.

Let B_i denote i 's utility for getting the promotion; by affine transformations of agents' utilities, if necessary, we may assume that $B_1 = B_2 = B$ and that the status-quo utility of no promotion is 0 for both agents. Let c_i denote the disutility incurred by agent i when he works in all T periods. We suppose that this disutility is uniform and additive across time. Thus $\frac{c_i}{T}$ is i 's disutility from work in any single period. For simplicity of exposition we also suppose $c_1 = c_2$ (see, however, Subsection 4.2).

To define the payoff of an agent i at any terminal node ω , consider the unique path in the tree from ω^* (the start of the game) to ω . Agent i 's disutility D_i^ω at ω is the sum of the one-period disutilities of his effort on this path. The probability p_i^ω that i gets promoted at ω is determined as follows. The principal samples n different periods from 1 to T , uniformly and randomly, and compares total outputs of agents 1 and 2 at the sampled periods along the path from ω^* to ω . An agent is promoted if his total output is higher than his rival's. In the event of a tie, each is promoted with probability $\frac{1}{2}$. Thus the expected payoff to agent i at the terminal node ω is:

$$p_i^\omega \cdot B - D_i^\omega.$$

This defines a T -period *promotion tournament* with sample size n , which we denote by $\Gamma_{p_1, p_2, c, B}(T, n)$ (or just $\Gamma(T, n)$). Notice that agents know the number n of days being sampled, but not which days; this is revealed to them only ex-post after the game is over.

Let $\Omega(T)$ denote the set of all nodes in the T -period game tree, at which agents move (i.e., all non-terminal nodes that do not correspond to chance moves). For any $\omega \in \Omega(T)$ we denote by $\Gamma_{p_1, p_2, c, B}(T, \omega, n)$ (or just $\Gamma(T, \omega, n)$) the subgame of $\Gamma_{p_1, p_2, c, B}(T, n)$ that starts at ω .

A (*behavioral*) *strategy* of an agent is given by a choice of probability distributions over his effort levels (work or shirk) at every node in $\Omega(T)$. We assume that the probabilities p_1, p_2 of production are common knowledge to the agents. Thus any pair of strategies of the two agents determines a probability distribution on the terminal nodes in the obvious manner, and thereby expected payoffs.

A pair of strategies in a game is called a *strategic equilibrium* (or SE) if no player can improve his payoff by unilaterally changing his own strategy. An SE (σ_1, σ_2) of $\Gamma(T, n)$ is called a (*subgame*) *perfect* SE (or PSE) if, for every $\omega \in \Omega(T)$, the restriction of strategies (σ_1, σ_2) to nodes of the subtree starting at ω constitute an SE of the subgame $\Gamma(T, \omega, n)$.

Remark 1 *A PSE always exists in a promotion tournament. Moreover, it exists even in strategies that are not conditioned on past effort levels, but only on past outcomes. We prove this in Section 5.*

3 The Main Results

Fix the parameters p_1, p_2, c, B , and w.l.o.g. assume that $p_1 \geq p_2$. Let $U : [0, 1]^2 \rightarrow R_+$ be a continuous non-decreasing function, and assume $U(x, y) = 0$ if and only if $x = 0$. The principal's objective in our world is to maximize $U\left(\frac{q}{T}, p\right)$, where q is the expected total output of the agents, and p is the probability that the "right guy" (i.e., agent 1) is promoted. It is implicit in this formulation that the principal does not know agents' skills, i.e., their probabilities p_i of being productive. He values their output, and may also be interested in the correct promotion, provided its implementation does not reduce output to zero. (Such a function U can arise, for a risk-neutral principal, as an expectation of utility defined on pure outcomes: if agent i is promoted and the total output is q , the principal obtains $a_i q$ utiles from this

outcome, where $a_1 \geq a_2$. When $a_1 = a_2$ the principal cares only about the total expected output.)

Let $\sigma = (\sigma_1, \sigma_2)$ be a pair of agents' strategies in $\Gamma(T, n)$. Denote $q(\sigma) \equiv$ the total expected output under σ , $p(\sigma) \equiv$ the probability that 1 is promoted under σ , and $U(\sigma) \equiv U\left(\frac{q(\sigma)}{T}, p(\sigma)\right)$. Define²

$$\bar{U}(T, n) = \max \{U(\sigma) \mid \sigma \text{ is a PSE of } \Gamma(T, n)\}$$

and

$$\underline{U}(T, n) = \min \{U(\sigma) \mid \sigma \text{ is a PSE of } \Gamma(T, n)\}.$$

We say that sample size n *dominates* sample size k in T -period promotion tournaments if $\underline{U}(n) > \bar{U}(k)$. Sample size k is said to be *undominated* if no sample size dominates it. Let

$$\eta(T) = \max \{k = 1, \dots, T \mid k \text{ is undominated}\}.$$

No matter how pessimistic and cautious the principal may be, he will reject any sample size $k > \eta(T)$ because the worst PSE outcome under $\eta(T)$ is still better for him than the best PSE outcome under k . Thus $\eta(T)$ serves as natural upper bound on acceptable sample sizes in T -period promotion tournaments.

We make the following assumptions on the parameters p_1, p_2, c and B :

(I) Agent 1 is sufficiently more skilled³ than 2:

$$p_1 > 3p_2.$$

(II) Agents' disutility from work is not too high relative to their productivity and the value they place on promotion:

$$\text{a) } c < \left(p_1(1 - p_2) - \frac{1}{2}\right) B;$$

$$\text{b) } c < p_2(1 - p_1) B.$$

Theorem 1 *Assume (I) and either (IIa) or (IIb). Then for fixed p_1, p_2, c, B*

$$\frac{\eta(T)}{T} \rightarrow 0$$

as $T \rightarrow \infty$.

²We write max, min instead of sup, inf because the set of PSE's is compact in every $\Gamma(T, n)$, and nonempty by Remark 1.

³We do not know if this condition can be dropped (see Subsection 4.3).

Theorem 1 is based upon the following three results, which are of interest in their own right.

Theorem 2 *Assume (I) and let $\varepsilon, s \in (0, 1]$. There exists $\bar{T} > 0$ such that, if T and n are integers with $T \geq \bar{T}$ and $T \geq n \geq sT$, then the expected number of workdays of each agent does not exceed εT in any PSE of the game $\Gamma_{p_1, p_2, c, B}(T, n)$.*

Theorem 2 reveals that large samples (which are non-vanishing fractions of T) are quite disastrous: the number of workdays, and hence expected output, becomes an arbitrarily small fraction of T for large enough T (choose ε small for fixed s).

The other two results (Propositions 3 and 4 below) show that, under sample size 1 or 2, the expected output at any PSE is asymptotically above a positive constant fraction of T . Thus, in conjunction with Theorem 2, they immediately yield Theorem 1: large samples (non-vanishing fractions of T) are dominated by samples even as small as 1.

Proposition 3 *Assume (IIa). Then I works at least $\frac{(p_1(1-p_2)-\frac{1}{2})^{B-c}}{B}T$ periods (in expectation) in every SE of $\Gamma_{p_1, p_2, c, B}(T, 1)$.*

Proposition 4 *Assume (IIb). Then both agents work in every period with probability 1 in any PSE of $\Gamma_{p_1, p_2, c, B}(T, 1)$.*

Under a stronger condition, Proposition 5 below guarantees that both agents work all the time in any PSE if the sample size is 2. Thus, so long as the principal cares about the correct promotion, sample size 2 dominates sample size 1 because the probability of detecting agent 1 as a winner goes up. This indicates that in general $\eta(T)$ exceeds 1.

Proposition 5 *Assume that*

$$c < \frac{p_2 B}{3} \cdot \min(p_1, 1 - p_1). \quad (1)$$

Then both agents work in every period with probability 1 in any PSE of $\Gamma_{p_1, p_2, c, B}(T, 2)$ for $T \geq 6$.

Proofs of all our results are given in Section 5.

4 Remarks

4.1 The rate of drop of sample size in Theorem 1

1. Theorem 2 shows that, given $0 < s \leq 1$, all sample sizes above sT make the total output fall below an arbitrarily small fraction of T , provided T is large enough. It can be shown, by mimicking the existing proof, that this upper bound on undominated sample sizes can be lowered to $s\sqrt{T}$, or in general any function $o(T)$ that has the following two properties:

(i) $o(T)$ is a vanishing fraction of T , i.e.,

$$\lim_{T \rightarrow \infty} \frac{o(T)}{T} = 0;$$

(ii) $o(T)$ grows much faster than $\ln T$, i.e.,

$$\lim_{T \rightarrow \infty} \frac{o(T)}{\ln T} = \infty.$$

Thus $\eta(T)$ is (asymptotically in T) bounded from above by any such function $o(T)$.

2. It can also be shown that, under the above assumption on the sample size, the expected total output in PSE is itself asymptotically bounded by a function $o'(T)$ with properties (i) and (ii).

4.2 Generalizations and Variations of the Model

1. The assumption of binary effort levels (work, shirk) and production levels (zero, one) can be replaced by more general assumptions, and certain versions of Theorems 1 and 2 will still hold. For instance, assume that there are finitely many agents, whose effort levels and outputs lie in finite sets, and whose disutilities from effort may differ but are positive whenever the effort leads to positive expected output. Denote by \bar{E}_i the maximal expected one-period output of agent i over all possible effort levels. It can be shown that there exists $d > 1$ which makes the following result true. Suppose there is an agent i with

$$\bar{E}_i > d\bar{E}_j$$

for all $j \neq i$. Then, given $\varepsilon, s \in (0, 1]$, there exists $\bar{T} > 0$ such that, for any integers T and n with $T \geq \bar{T}$ and $T \geq n \geq sT$, the expected number of

workdays of each agent does not exceed εT in any PSE of the promotion tournament of length T with sample size n . The proof of this result follows closely the one given in the paper, but is messier.

2. Suppose the agents move sequentially rather than simultaneously (think of day and night shifts in a factory), and have perfect information in the game. In this case Kuhn’s theorem guarantees the existence of PSE in pure strategies. Our main results are affected only slightly. Suppose that the periods in which agents 1 and 2 move have a “ δ -uniform” mix, for some sufficiently small δ ; i.e., there exists k independent of T such that in any sequence of k periods each agent moves at least $\left(\frac{1}{2} - \delta\right) k$ times. Then there exists $d(\delta) > 1$ such that the assertions of Theorems 1 and 2 hold, provided assumption (I) is replaced by:

$$p_1 > d(\delta) p_2.$$

In fact, we may also allow any mix of simultaneous and sequential moves in the above scenario.

When the game has only sequential moves, undominated sample sizes are generically optimal. If we independently perturb, at every node of the game tree, both the disutilities of work and the probabilities of production (in small neighborhoods of c and p_i), our main results remain intact. Generically in the perturbation, however, PSE is unique and consists of pure strategies for all sample sizes. And then undominated sample sizes are just those that maximize the principal’s utility.

4.3 Open Questions

1. Do Theorems 1 and 2 hold if assumption (I) is weakened to $p_1 > p_2$? If this were the case, we could probably extend our analysis to general tournaments and get similar results. In a promotion tournament there is only one prize, but in a general tournament there are multiple prizes, awarded to agents based on the ranking of their outputs. Our current analysis as it stands also extends to general tournaments, but under progressively restrictive assumptions: agent 1 must be much stronger than agent 2, 2 much stronger than 3, etc.

2. What happens when $p_1 = p_2$: is there $\varepsilon_0 > 0$ such that the expected number of workdays of the agents exceeds $\varepsilon_0 T$, for any number of periods T and any sample size?

3. Do Theorems 1 and 2 hold if we require a winning player in $\Gamma(T, n)$ to be ahead of his rival by αT , for some $\alpha > 0$?

4. Suppose that the principal does not care about the correct promotion, but only about total output, i.e., U is a function of only the first coordinate. It is easy to see that if the sample size $n \geq 3$ then, for all sufficiently large T , there is no PSE in which agents work all the time with probability 1. Then Propositions 4 and 5 imply that the only undominated sample sizes are 1 and 2, provided (1) holds. Is this conclusion still true without assumption (1), i.e., what happens when no sample size is capable of inducing full-time work?

5 Proofs

5.1 Lemmas

Let N denote the set of natural numbers, and assume (I) throughout this subsection.

Lemma 6 *Let $0 < \delta$. Given $D \in N$ and a set of D periods, suppose that agent i works unconditionally⁴ in each of these periods. Then the probability that he produces more than $(p_i + \delta)D$ units of output during these periods, and the probability that he produces less than $(p_i - \delta)D$ units, are both at most e^{-aD} , for some $a > 0$ independent of D .*

Proof. The proof is trivial if $p_i = 0$ or 1 , so we assume $0 < p_i < 1$. The number of units S_D that agent i produces in the D periods is a sum of D independent and identically distributed random variables X_t , such that $X_t = 1$ if i produces one unit of output in period t (which happens with probability p_i), and $X_t = 0$ otherwise. The variable S_D has mean $p_i D$ and standard deviation $\sqrt{p_i(1-p_i)D}$. The probability that i produces more than $(p_i + \delta)D$ units is therefore

⁴If an agent works at every node of a subset Ω of $\Omega(T)$, we say that he works unconditionally in Ω . If Ω is the set of all nodes in period t , we say that the agent works unconditionally in period t . Unconditional shirking is defined similarly.

$$\Pr(S_D - p_i D \geq \delta D) = \Pr\left(\frac{S_D}{\sqrt{p_i(1-p_i)D}} - \frac{p_i D}{\sqrt{p_i(1-p_i)D}} \geq \frac{\delta\sqrt{D}}{\sqrt{p_i(1-p_i)}}\right). \quad (2)$$

Denoting

$$d_i = \min\left(\frac{\delta}{\sqrt{p_i(1-p_i)}}, \sqrt{p_i(1-p_i)}\right),$$

the probability in (2) is at most

$$\Pr\left(\frac{S_D}{\sqrt{p_i(1-p_i)D}} - \frac{p_i D}{\sqrt{p_i(1-p_i)D}} \geq d_i\sqrt{D}\right).$$

By A (i) on p. 254 of Loeve (1960), this probability is at most

$$\exp\left(-\frac{(d_i\sqrt{D})^2}{2}\left(1 - \frac{d_i\sqrt{D}}{2\sqrt{p_i(1-p_i)D}}\right)\right) \leq \exp\left(-\frac{d_i^2 D}{4}\right).$$

It can be shown similarly that $\exp\left(-\frac{d_i^2 D}{4}\right)$ bounds $\Pr(S_D - p_i D \leq -\delta D)$ from above. Taking $a = \min_{i=1,2}\left(\frac{d_i^2}{4}\right)$ finishes the proof. ■

Lemma 7 *Let $0 < \delta \leq p_1 - p_2$. Given $D \in \mathbb{N}$ and a set of D periods, suppose that agent 1 works unconditionally in each of these periods. The probability that he produces at least $(p_1 - p_2 - \delta)D$ units more than 2 during these periods is bounded from below by $1 - 2e^{-aD}$ for some $a > 0$ independent of D .*

Proof. The required probability is minimal when 2 also works unconditionally in each of the D periods. It is bounded from below by the probability that 1 produces at least $(p_1 - \frac{\delta}{2})D$ units, and 2 produces at most $(p_2 + \frac{\delta}{2})D$ units during these periods. By independence of production, this is the product of the probabilities of the two events. By Lemma 6, the required probability is at least $(1 - e^{-aD})^2 \geq 1 - 2e^{-aD}$, for $a > 0$ independent of D . ■

Given a node ω , denote the total output of i on the path from ω^* to ω by $A_i(\omega)$, for $i \in \{1, 2\}$.

Lemma 8 *Let $\varepsilon, s \in (0, 1]$. There exists $a > 0$ such that, for every T and n in N with $T \geq \max(n, 2) \geq n \geq sT$, and every terminal node ω of the game $\Gamma(T, n)$ with*

$$A_1(\omega) \geq A_2(\omega) + \varepsilon T,$$

the probability that 1 wins the promotion is at least $1 - 2e^{-aT}$, conditional on the realization of ω .

Proof. Denote by S^i the number of i 's outputs sampled by the principal on the path from ω^* to ω ; note that the expectation of S^i is equal to $\frac{n}{T}A_i(\omega)$. The probabilities below are all conditional on the realization of ω :

$$\Pr(1 \text{ wins the promotion}) \geq \Pr\left(S^1 - \frac{n}{T}A_1(\omega) \geq -\frac{\varepsilon n}{4} \text{ and } S^2 - \frac{n}{T}A_2(\omega) \leq \frac{\varepsilon n}{4}\right) \quad (3)$$

$$\geq 1 - \Pr\left(S^1 - \frac{n}{T}A_1(\omega) < -\frac{\varepsilon n}{4}\right) - \Pr\left(S^2 - \frac{n}{T}A_2(\omega) > \frac{\varepsilon n}{4}\right). \quad (4)$$

Let us estimate the second term in (4). Suppose first that $A_1(\omega) \leq \left(1 - \frac{\varepsilon n}{4T}\right)T$. For an integer A between εT and $\left(1 - \frac{\varepsilon n}{4T}\right)T$ define

$$\sigma_A = \sqrt{\frac{1}{T-1} \left(A \left(1 - \frac{A}{T}\right)^2 + (T-A) \left(\frac{A}{T}\right)^2 \right)}.$$

Also denote

$$\bar{\sigma} = \max_{T \geq 2, \varepsilon T \leq A \leq \left(1 - \frac{\varepsilon n}{4}\right)T} \sigma_A,$$

$$\underline{\sigma} = \min_{T \geq 2, \varepsilon T \leq A \leq \left(1 - \frac{\varepsilon n}{4}\right)T} \sigma_A,$$

$$\varepsilon'' = \frac{n}{T} \min\left(\frac{\varepsilon}{4}, \underline{\sigma}^2\right),$$

and note that $\infty > \bar{\sigma} > \underline{\sigma} > 0$. Now,

$$\begin{aligned} \Pr\left(S^1 - \frac{n}{T}A_1(\omega) < -\frac{\varepsilon n}{4}\right) &\leq \Pr\left(S^1 - \frac{n}{T}A_1(\omega) < -\varepsilon''T\right) \\ &= \Pr\left(\frac{S^1}{\sigma_{A_1(\omega)}\sqrt{n}} - \frac{\frac{n}{T}A_1(\omega)}{\sigma_{A_1(\omega)}\sqrt{n}} < -\frac{\varepsilon''T}{\sigma_{A_1(\omega)}\sqrt{n}}\right). \end{aligned}$$

By Lemma 6.3 of Rosen (1965) this probability is at most

$$\exp \left(-\frac{\left(\frac{\varepsilon'' T}{\sigma_{A_1(\omega)} \sqrt{n}} \right)^2}{2} \left(1 - \frac{\frac{\varepsilon'' T}{\sigma_{A_1(\omega)} \sqrt{n}}}{2\sigma_{A_1(\omega)} \sqrt{n}} \right) \right).$$

This expression equals

$$\begin{aligned} \exp \left(-\frac{(\varepsilon'')^2 T}{2\sigma_{A_1(\omega)}^2 n} \left(1 - \frac{\varepsilon''}{2\sigma_{A_1(\omega)} \frac{n}{T}} \right) \right) &\leq \exp \left(-\frac{(\varepsilon'')^2 T}{4\sigma_{A_1(\omega)}^2 \frac{n}{T}} \right) \leq \exp \left(-\frac{(\varepsilon'')^2 T}{4\bar{\sigma}^2 \frac{n}{T}} \right) \\ &= \exp \left(-\frac{\left(\frac{n}{T} \min \left(\frac{\varepsilon}{4}, \sigma^2 \right) \right)^2 T}{4\bar{\sigma}^2 \frac{n}{T}} \right) \equiv \exp(-na_1). \end{aligned}$$

Therefore

$$\Pr \left(S^1 - \frac{n}{T} A_1(\omega) < -\frac{\varepsilon n}{4} \right) \leq \exp(-na_1).$$

Now, if $A_1(\omega) > \left(1 - \frac{\varepsilon n}{4T}\right) T$, then the principle samples at most $\frac{\varepsilon n}{4}$ periods on which 1 did not produce, and so at least $n - \frac{\varepsilon n}{4} \left(\geq \frac{n}{T} A_1(\omega) - \frac{\varepsilon n}{4}\right)$ units of 1's output are sampled. Therefore $\Pr \left(S^1 - \frac{n}{T} A_1(\omega) < -\frac{\varepsilon n}{4} \right) = 0$, and, in particular, is below $\exp(-a_1 n)$.

The existence of $a_2 > 0$ such that

$$\Pr \left(S^2 - \frac{n}{T} A_2(\omega) > \frac{\varepsilon n}{4} \right) \leq \exp(-a_2 n)$$

for all T is shown similarly. It follows from (3) and (4) that

$$\begin{aligned} \Pr(I \text{ wins the promotion}) &\geq 1 - 2 \exp(-\min(a_1, a_2) n) \\ &\geq 1 - 2 \exp(-\min(a_1, a_2) sT) \equiv 1 - 2 \exp(-aT). \end{aligned}$$

■

Lemma 9 *Assume (I), and let $\varepsilon, s \in (0, 1]$. There exists $\bar{T} > 0$ with the following property: suppose that T and n are integers such that $T \geq \bar{T}$ and $T \geq n \geq sT$, (σ_1, σ_2) is a PSE of $\Gamma_{p_1, p_2, c, B}(T, n)$, and $\omega \in \Omega(T)$ is such that*

$$A_1(\omega) \geq A_2(\omega) + \varepsilon T; \tag{5}$$

then the strategies (σ_1, σ_2) , restricted to the subgame $\Gamma(T, \omega, n)$, induce a unique path along which both agents shirk. In other words, if agents employ (σ_1, σ_2) in $\Gamma(T, \omega, n)$, they shirk at every period of the game with probability 1.

Proof. Set $\varepsilon' = \frac{\varepsilon}{3}$ and $\delta = \frac{p_1 - 3p_2}{4}$. Let $\bar{T} > \frac{3}{\varepsilon}$ be such that for any $T \geq \bar{T}$ the following inequalities hold for positive a_1, a_2, a_3 that will be specified later:

(a)

$$\frac{c}{T} > 2B \left(e^{-a_1 T} + e^{-a_3 T} \right);$$

(b)

$$\begin{aligned} 2 \left(1 - \left(1 - 2e^{-a_2 T} \right) \left(1 - 2e^{-a_1 T} \right) \right) B + \left(1 - 2e^{-a_2 T} \right) \frac{2\varepsilon' T (p_2 + \delta)}{p_1 - p_2 - \delta} c + 4e^{-a_2 T} c - \frac{[\varepsilon' T] c}{T} \\ < 2e^{-a_1 T} B - \frac{c}{T}; \end{aligned}$$

(c)

$$\left[\frac{\varepsilon' (p_2 + \delta)}{p_1 - p_2 - \delta} T \right] (p_1 - p_2 - \delta) > \left(p_2 + \frac{\delta}{2} \right) \varepsilon' T,$$

where $[m]$ stands for the integer part of the real number m . Also select and fix a PSE in every promotion tournament $\Gamma(T, n)$, for all T and $T \geq n \geq sT$.

Given integers T and n such that $T \geq \bar{T}$ and $T \geq n \geq sT$, we prove the lemma for every $\omega \in \Omega(T)$ satisfying (5), by backward induction on $A_1(\omega)$. Fix one such ω . Since only the subtree starting at ω will be considered, we can add to all payoffs of agent i the number $C_i(\omega)$ - his disutility from work incurred prior to ω . The newly defined payoffs ignore costs from work prior to ω , and do not change strategic behavior in the subtree.

First suppose that there are at most $[\varepsilon' T]$ periods from ω to the end of the game (which must be the case if $A_1(\omega) \geq (1 - \varepsilon')T$). Let a_1 be the positive constant from Lemma 8, applied for ε' and s . Then the probability that 1 wins the promotion, conditional on the PSE path reaching ω , is at least $1 - 2e^{-a_1 T}$, no matter what the agents do after ω (indeed, the difference $A_1(\omega') - A_2(\omega')$ is at least $\varepsilon' T$ for any terminal node ω' following ω). Thus if agent 1 chooses to shirk at ω and then follow his PSE strategy, he gets at least $(1 - 2e^{-a_1 T})B$ (this payoff can be guaranteed by doing nothing). If agent 1 works at ω and then follows his PSE strategy, he incurs a disutility of $\frac{c}{T}$ and can get at most B . The change in utility is thus at most $2Be^{-a_1 T} - \frac{c}{T}$.

Since $T \geq \bar{T}$, (a) in the definition of \bar{T} implies that it is always better to shirk at ω , and so 1's PSE strategy will not put a positive probability on work at ω . Similarly, if 2 shirks at ω and then follows his PSE strategy, he receives some positive payoff (which can be guaranteed by doing nothing). If, however, 2 works at ω and then follows his PSE strategy, he can get at most $2e^{-a_1 T} B$, and incurs a cost of $\frac{\epsilon}{T}$. The change in utility is again negative by (a), hence 2's PSE strategy will tell him to shirk at ω . A repetition of these arguments shows that no agent works at any ω' following ω , and so there is no work after ω .

Next suppose that there are more than $[\epsilon' T]$ periods from ω to the end of the game (and so $A_1(\omega) < (1 - \epsilon') T$), and the assertion of the lemma holds for all ω' with $A_1(\omega') > A_1(\omega)$. We will show that 2's PSE strategy tells him to shirk at ω .

We first prove that if 2 decides to work at ω and to follow his PSE strategy thereafter, and 1 sticks to his PSE strategy at all times, then 2's expected payoff is negative. Denote by Ω' the set of nodes $\omega' \in \Omega(T)$ subject to the following three conditions:

- i) 2 works at ω' ;
- ii) there are exactly $[\epsilon' T]$ periods on the path leading from ω to ω' on which 2 works;
- iii) there are more than $[\epsilon' T]$ periods from ω' to the end of the game.

Let Ω'' be the set of nodes $\omega' \in \Omega'$ such that 2 produces at most $(p_2 + \frac{\delta}{2}) \epsilon' T$ units of output during the $[\epsilon' T]$ workdays between ω and ω' . Denote by F the event that a node in Ω' is reached, and by $E \subset F$ the event that a node in Ω'' is reached. Conditional on E and ω' being a node reached in Ω'' , we come to the game $\Gamma(T, \omega', n)$. Denote by $G_1(\omega')$ the PSE payoff of agent 1 in this game, which ignores costs of work incurred prior to ω' .

To bound $G_1(\omega')$ from below we consider the following strategy of 1 (that may differ from his PSE strategy): let 1 work unconditionally for $\lceil \frac{\epsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} T \rceil$ consecutive periods after ω' (until, say, $\omega'' \in \Omega(T)$), and then switch to his PSE strategy. By Lemma 7 there is $a_2 > 0$ independent of T such that, in the $\lceil \frac{\epsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} T \rceil$ periods following ω' , 1 produces at least $\lceil \frac{\epsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} T \rceil (p_1 - p_2 - \delta)$ units more than 2 does during these periods, with probability at least $1 - 2e^{-a_2 T}$. Since $T \geq \bar{T}$, (c) in the definition of \bar{T} implies that the probability of $A_1(\omega'') > A_2(\omega'') + \epsilon T$ and $A_1(\omega'') > A_1(\omega)$ is at least $1 - 2e^{-a_2 T}$. By our induction assumption neither agent works after ω'' on the path induced by the PSE strategies, with probability at least $1 - 2e^{-a_2 T}$. Given that the

agents do not work after ω'' on this path and $A_1(\omega'') > A_2(\omega'') + \varepsilon T$, the probability that 1 wins the promotion is at least $1 - 2e^{-a_1 T}$, by the definition of a_1 . Thus

$$G_1(\omega') \geq (1 - 2e^{-a_2 T}) \left((1 - 2e^{-a_1 T}) B - \frac{\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c \right) - 2e^{-a_2 T} \cdot c. \quad (6)$$

Using (6), the probability α that 1 wins the promotion in the game $\Gamma(T, \omega', n)$ can be estimated. Let β be the probability of a tie. Then we must have

$$\alpha B + \beta \frac{1}{2} B \geq G_1(\omega')$$

$$\geq (1 - 2e^{-a_2 T}) \left((1 - 2e^{-a_1 T}) B - (1 - 2e^{-a_2 T}) \frac{\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c - 2e^{-a_2 T} c \right),$$

and so (taking the worst case: $\beta = 1 - \alpha$)

$$\alpha \geq \frac{\left((1 - 2e^{-a_2 T}) \left((1 - 2e^{-a_1 T}) - \frac{1}{2} \right) B - (1 - 2e^{-a_2 T}) \frac{\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c - 2e^{-a_2 T} c \right)}{\frac{1}{2} B}.$$

Therefore 2's probability of getting the promotion (through a win or a tie) is

$$1 - \alpha \leq \frac{\left(1 - \left(1 - 2e^{-a_2 T} \right) \left(1 - 2e^{-a_1 T} \right) \right) B + \left(1 - 2e^{-a_2 T} \right) \frac{\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c + 2e^{-a_2 T} c}{\frac{1}{2} B}.$$

Thus the PSE payoff of 2 in the game $\Gamma(T, \omega', n)$ is at most

$$(1 - \alpha) B$$

$$= 2 \left(1 - \left(1 - 2e^{-a_2 T} \right) \left(1 - 2e^{-a_1 T} \right) \right) B + \left(1 - 2e^{-a_2 T} \right) \frac{2\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c + 4e^{-a_2 T} c$$

(ignoring costs incurred from work prior to ω'). If the costs of work are ignored prior to ω only, 2's payoff is at most

$$2 \left(1 - \left(1 - 2e^{-a_2 T} \right) \left(1 - 2e^{-a_1 T} \right) \right) B + \left(1 - 2e^{-a_2 T} \right) \frac{2\varepsilon'(p_2 + \delta)}{p_1 - p_2 - \delta} c + 4e^{-a_2 T} c - \frac{[\varepsilon' T] c}{T}.$$

We now proceed to estimate the probability of $F - E$. Consider a sequence of independent and identically distributed random variables $(Y_n)_{n=1}^{\infty}$, such

that $Y_n = 1$ with probability p_2 and $Y_n = 0$ with probability $1 - p_2$. Let us count only those periods, starting from ω , on which agent 2 works. Think of Y_n as the outcome of 2's work at period n (of course only a finite number of Y_n 's are ever looked at). If the event $F - E$ occurs, then $\frac{\sum_{n=1}^{\lceil \varepsilon' T \rceil} Y_n}{\lceil \varepsilon' T \rceil} > \left(p_2 + \frac{\delta}{2}\right)$. This inequality holds with probability at most $2e^{-a_3 T}$, for some $a_3 > 0$ which is independent of T - this is shown by the same arguments as in Lemma 6. Therefore the probability of $F - E$ is bounded from above by $2e^{-a_3 T}$.

Conditional on F^c (i.e., F does not occur) it is clear that

$$A_1(\omega) \geq A_2(\omega) + (\varepsilon - 2\varepsilon')T = A_2(\omega) + \varepsilon'T,$$

and so the probability that 2 wins or ties is at most $2e^{-a_1 T}$, by the definition of a_1 . Therefore 2's payoff is at most $2e^{-a_1 T}B - \frac{c}{T}$ (here we use our assumption that 2 works at ω , and thus incurs a disutility of at least $\frac{c}{T}$).

Therefore the payoff that 2 receives is bounded from above by

$$\begin{aligned} & \left(2 \left(1 - \left(1 - 2e^{-a_2 T}\right) \left(1 - 2e^{-a_1 T}\right)\right) B + \left(1 - 2e^{-a_2 T}\right) \frac{2\varepsilon'T(p_2 + \delta)}{p_1 - p_2 - \delta} c + 4e^{-a_2 T} c - \frac{\lceil \varepsilon'T \rceil c}{T}\right) \\ & \times \Pr(E) + B \Pr(F - E) + \left(2e^{-a_1 T} B - \frac{c}{T}\right) \Pr(F^c). \end{aligned}$$

Using (b) in the definition of \bar{T} , the above expression is maximal when $\Pr(E) = 0$, and so this expression is bounded from above by

$$B \Pr(F - E) + \left(2e^{-a_4 T} B - \frac{c}{T}\right) \leq 2e^{-a_1 T} B + 2e^{-a_3 T} B - \frac{c}{T}.$$

By (a), this number is negative. Therefore 2 receives a negative payoff if he works at ω and follows his PSE strategy thereafter. On the other hand, if 2 shirks at all periods following ω , he receives a non-negative payoff. This implies that if 2 shirks at ω and then follows his PSE strategy, he also receives a non-negative payoff. We deduce that it is better for 2 to shirk at ω and to follow his PSE strategy thereafter, than to work at ω and to follow his PSE strategy thereafter, provided 1 sticks to his PSE strategy. It follows that 2's PSE strategy instructs him to shirk at ω with probability 1.

Now, let 1 unconditionally shirk, starting at ω . We know that 2's PSE strategy tells him to shirk at ω . This leads to a later node ω' , in which $A_1(\omega') \geq A_2(\omega') + \varepsilon T$. By the same proof 2 also shirks at ω' . Successive repetitions of this argument show that if 1 unconditionally shirks, starting

at ω , so does 2. This gives 1 a payoff of at least $(1 - 2e^{-a_1 T})B$. Thus, if 1 shirks at ω and follows his PSE strategy thereafter, he is also getting at least $(1 - 2e^{-a_1 T})B$. On the other hand, if 1 works at ω and follows his PSE thereafter, he gets at most $B - \frac{c}{T}$. Using (a), this is less than $(1 - 2e^{-a_1 T})B$. Thus the PSE strategy of 1 will instruct him to shirk at ω .

We have shown that both agents shirk at ω in PSE. This leads to a later node, for which the same proof works: PSE strategies tell agents to shirk there. Repetitive applications of this argument establish that neither agent works on the PSE path starting at ω . This completes the inductive step, and concludes the proof. ■

5.2 Proof of Theorem 2

Fix a PSE in every promotion tournament $\Gamma(T, n)$, for all T and $T \geq n \geq sT$. Let $\varepsilon' = \frac{\varepsilon}{3}$. For a_1, a_2 , and T' that will be specified later, we take $\bar{T} > 0$ be an integer such that for all integers $T \geq \bar{T}$:

1)

$$(1 - 2e^{-a_1 T}) \left[(1 - 2e^{-a_2 T}) B - \varepsilon' c \right] - 2e^{-a_1 T} c > B - \varepsilon c;$$

2)

$$2 \left(1 - (1 - 2e^{-a_1 T}) (1 - 2e^{-a_2 T}) \right) B + 2 (1 - 2e^{-a_1 T}) \varepsilon' c + 4e^{-a_1 T} c < \varepsilon c;$$

3)

$$T \geq T'.$$

Assume that 2 sticks to his PSE strategy at all times, and consider the following strategy of 1 in $\Gamma(T, n)$: let 1 work for the first $\lceil \varepsilon' T \rceil$ periods, and then switch to his PSE strategy. With probability at least $1 - 2e^{-a_1 T}$ he will reach a node in which his output is at least $\frac{p_1 - p_2}{2} \cdot \lceil \varepsilon' T \rceil$ units ahead of 2 (such $a_1 > 0$ exists independently of T and $n \geq sT$, by Lemma 7). For some $T' = T'(s)$ and every $T \geq T'$ and $n \geq sT$, the agents will shirk starting at this node (Lemma 9), in which case the probability of 1 winning the promotion is at least $1 - 2e^{-a_2 T}$ (for some $a_2 > 0$ independent of T and $n \geq sT$, by Lemma 8). Therefore 1 can guarantee a payoff of at least $(1 - 2e^{-a_1 T}) \left((1 - 2e^{-a_2 T}) B - \varepsilon' c \right) - 2e^{-a_1 T} c$. If, however, the expected number of workdays of 1 in the PSE exceeds εT , then 1 is getting at

most $B - \varepsilon c$, which is below the guaranteed payoff if $T \geq \bar{T}$. So 1 does not work more than εT periods in expectation.

Employing an argument similar to the one used in Lemma 9 for the computation of α (the probability of 1 winning the promotion) one can show that, given that 1 can guarantee a payoff of at least $(1 - 2e^{-a_1 T}) \left((1 - 2e^{-a_2 T}) B - \varepsilon' c \right) - 2e^{-a_1 T} c$, the probability of 2 getting the promotion (through a win or a tie) is at most

$$\frac{\left(1 - \left(1 - 2e^{-a_1 T}\right) \left(1 - 2e^{-a_2 T}\right)\right) B + \left(1 - 2e^{-a_1 T}\right) \varepsilon' c + 2e^{-a_1 T} c}{\frac{1}{2} B}.$$

So, if the expected number of 2's workdays exceeds εT , then 2's payoff in the PSE is at most

$$2 \left(1 - \left(1 - 2e^{-a_1 T}\right) \left(1 - 2e^{-a_2 T}\right)\right) B + 2 \left(1 - 2e^{-a_1 T}\right) \varepsilon' c + 4e^{-a_1 T} c - \varepsilon c.$$

Therefore, if $T \geq \bar{T}$, 2 would be better off shirking, contradicting the definition of PSE. This finishes the proof.

5.3 Proof of Theorem 1

If (IIa) or (IIb) are assumed, then Propositions 3 and 4 guarantee an existence of $\delta > 0$ such that $\frac{q(\sigma)}{T} \geq \delta$ at any PSE σ of $\Gamma(T, 1)$. Thus $\underline{U}(T, 1) > 0$.

Since $\lim_{\varepsilon \rightarrow 0} U(\varepsilon, 1) = 0$ by the continuity of U , we can find $\varepsilon > 0$ such that $U(\varepsilon, 1) < \underline{U}(T, 1)$. Theorem 2 implies that for any $1 \geq s > 0$ there exists $\bar{T} = \bar{T}(\varepsilon, s) > 0$ such that, for all $T > \bar{T}$ and $T \geq n \geq sT$, and a PSE σ of $\Gamma(T, n)$,

$$\bar{q}(\sigma) \leq \varepsilon T.$$

Since U is non-decreasing,

$$\bar{U}(T, n) < U(\varepsilon, 1),$$

and so $\bar{U}(T, n) < \underline{U}(T, 1)$. Sample size 1 thus dominates all sample sizes greater than sT in T -period promotion tournaments, for $T > \bar{T}$. Therefore

$$\limsup_{T \rightarrow \infty} \frac{\eta(T)}{T} \leq s.$$

Since the inequality holds for any $1 \geq s > 0$, we deduce that

$$\lim_{T \rightarrow \infty} \frac{\eta(T)}{T} = 0.$$

5.4 Proof of Proposition 3

If 1 works unconditionally in all periods, he guarantees a payoff of at least $-c + p_1(1 - p_2)B$ against any strategy of 2 (this is the actual payoff in the worst case scenario from 1's perspective: 2 also works unconditionally). If the expected number of workdays of 1 is at most εT , for $0 < \varepsilon \leq 1$, he gets at most $\varepsilon B + \frac{1}{2}B$ (assuming the best case scenario from 1's perspective: 2 shirks unconditionally, and 1 succeeds to produce whenever he works). For every $0 < \varepsilon < \frac{(p_1(1-p_2) - \frac{1}{2})B - c}{B}$

$$\varepsilon B + \frac{1}{2}B < -c + p_1(1 - p_2)B.$$

This means that, if the expected number of 1's workdays in any SE is less than εT , then 1 does not get the minimum he can guarantee, contradicting the definition of SE.

5.5 Proof of Proposition 4

Proof. We will show by backward induction that in every period both agents work unconditionally in any PSE of $\Gamma(T, 1)$. Let $\omega \in \Omega(T)$ be a node in the last period T . Then, by choosing to work rather than shirk at ω , agent i increases his probability to win by at least $\frac{1}{T}p_i(1 - p_j)$, whatever the action of the other agent j may be at ω ($\frac{1}{T}$ is the probability that the last period is sampled; $p_i(1 - p_j)$ bounds from below the probability that i produces one unit of output at this period, whereas his rival j does not). It is thus clear that the change in i 's payoff when he switches from shirk to work at ω is at least $\frac{1}{T}p_i(1 - p_j)B - \frac{c}{T}$, and this number is positive by assumption (IIb). Thus agents' PSE strategies will tell them to unconditionally work at the last period.

Now, assume that $\omega \in \Omega(T)$ is a node in period $t < T$, and that the agents work at every node following ω . No matter what agents' actions at ω may be, on subsequent periods both agents work by the inductive assumption. Therefore, as in the above paragraph, the change in payoff when agent i switches from shirk to work at ω is at least $\frac{1}{T}p_i(1 - p_j)B - \frac{c}{T}$, which is positive. Therefore the PSE strategies of both agents tell them to unconditionally work at ω . This finishes the inductive proof.

5.6 Proof of Proposition 5

The proof proceeds by backward induction, in a manner similar to the proof of Proposition 4 above. Let $\omega \in \Omega(T)$ be a node in the last period T , and let i be an agent whose total output along the path leading from ω^* to ω is no more than his rival's.

We claim that the change in i 's payoff, when he switches from shirk to work at ω , is positive, no matter what action the other agent j chooses at ω . Indeed, consider the event E that the last period is sampled, together with another period in which i does not produce more than the other agent j ; its probability is bounded from below by $2\left(\frac{1}{T} \cdot \frac{1}{3}\right)$ since $T \geq 6$. (In what follows the disutility from work will be ignored.) Given E and that both agents produce the same amount on the other sampled period, if j shirks at ω then, by working at ω , i ensures the payoff of $p_i B + (1 - p_i) \frac{B}{2}$, compared to $\frac{B}{2}$ when he shirks. If j works at ω then, by working at ω , i ensures the payoff of at least $p_i(1 - p_j)B + (1 - p_i)(1 - p_j)\frac{B}{2}$, compared to $(1 - p_j)\frac{B}{2}$ when he shirks. Next, given E and that agents produce different amounts on the other sampled period (i.e., j produces 1 and i produces 0), if j shirks at ω then, by working at ω , i ensures the payoff of $p_i \frac{B}{2}$, compared to 0 when he shirks. If j works at ω then, by working at ω , i ensures the payoff of at least $p_i(1 - p_j)\frac{B}{2}$, compared to 0 when he shirks. To summarize: conditional on E , the switch from work to shirk leads to a gain in payoff of at least $p_i(1 - p_j)\frac{B}{2}$. Finally observe that if E does not occur, the probability that i gets a promotion does not decrease when he switches at ω from shirk to work. 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 ω is at least

$$2\left(\frac{1}{T} \cdot \frac{1}{3}\right) p_i(1 - p_j) \frac{B}{2} - \frac{c}{T} = \frac{1}{T} p_i(1 - p_j) \frac{B}{3} - \frac{c}{T},$$

whatever is the action of j at ω may be. This number is positive by assumption (1). Therefore i 's PSE strategy will tell him to work at ω with probability 1.

Given that decision by i , the other agent j will also work at ω with probability 1. Indeed, it can be shown as before that the change in j 's payoff (ignoring the disutility from work) if he switches to work at ω is at least $p_j(1 - p_i)\frac{B}{2}$, provided the last period is sampled and that on the other sampled period j does not produce more than i . If the last period is sampled and on the other sampled period j produces 1 and i produces 0, then, by working

at ω , j can ensure the payoff of $p_j B + (1 - p_j) \left((1 - p_i) B + p_i \frac{B}{2} \right)$, compared to $(1 - p_i) B + p_i \frac{B}{2}$ when he shirks. Thus the change in j 's payoff resulting from the switch to work at ω is at least $p_j \frac{B}{2} \min(p_i, 1 - p_i)$, provided the last day is sampled. Thus the unconditional payoff of j (now taking the disutility from work into account) grows by

$$2 \frac{1}{T} p_j \min(p_i, 1 - p_i) \frac{B}{2} - \frac{c}{T} = \frac{1}{T} p_j \min(p_i, 1 - p_i) B - \frac{c}{T}$$

if j works, rather than shirks, at ω , and this number is positive, as can be seen from (1). Thus the PSE strategy of j tells him to work unconditionally at ω .

The proof of the inductive step is very close to the one just given above for the last period, and we omit it (all comparisons must be done now conditional on a path of outputs from the next period on, whose distribution is determined by the inductive assumption that agents work unconditionally in the sequel).

5.7 Proof of Remark 1

We construct a PSE (σ_1, σ_2) by (backward) induction on the time period to which a node in $\Omega(T)$ correspond. Partition first the nodes of the last period of the game (T -nodes) into *types*: two T -nodes are said to be of the same type if the paths leading from ω^* to them have the same history of outputs. At every T -node each agent has two actions (work, shirk), and a pair of actions, once taken, leads to a payoff. Thus at each T -node we have a one-shot two-person game, which possesses a mixed strategy equilibrium. Let (τ_1, τ_2) be one such equilibrium, and let σ_i tell agent i to employ τ_i upon reaching this node. Since this one-shot game is identical at all T -nodes of the same type, so are the sets of equilibria, and we can make our selection of (τ_1, τ_2) invariant at all nodes of the same type. Thus the strategies σ_i are “type-symmetric” in period T , i.e., prescribe identical behavior at all T -nodes of the same type.

Suppose now that type-symmetric strategies (σ_1, σ_2) have been constructed for all periods greater than t . Once again partition t -nodes into types based on the past history of outputs. At any t -node a one shot two-person game is defined, in which agents decide simultaneously to either work or shirk, and the payoff is determined based on the t -period outputs, and on the assumption that future behavior is according to the type-symmetric (σ_1, σ_2) .

A mixed strategy equilibrium exists in this game, and we pick one: $(\bar{\tau}_1, \bar{\tau}_2)$. Note that, at any two t -nodes of the same type, the same pair of actions leads to identical expected payoffs. This is so because (σ_1, σ_2) , being type-symmetric, does not distinguish between later nodes when they are preceded by the same history of outputs. Therefore we can make the choice of $(\bar{\tau}_1, \bar{\tau}_2)$ invariant at all t -nodes of the same type. As before, let σ_i tell agent i to employ $\bar{\tau}_i$ upon reaching a t -node. The strategies σ_i are clearly type-symmetric in period t .

The construction is finished when we get to the root of the game tree. The pair (σ_1, σ_2) is a PSE, by standard backward induction arguments (e.g., Kuhn (1953)).

REFERENCES

1. Cowen, T., and Glazer, A. (1996). "More Monitoring Can Induce Less Effort," *Journal of Economic Behavior and Organization* **30**, 113-23.
2. Dubey, P. and C. Wu (2000). "Competitive Prizes: When Less Scrutiny Induces More Effort," Cowles Foundation DP#1255, Yale University.
3. Green, J. R., and Stokey, N.L. (1983). "A Comparison of Tournaments and Contracts," *Journal of Political Economy* **91**, 349-64.
4. Kuhn, H. W. (1953). "Extensive Games and Problems of Information," in H. W. Kuhn and A. W. Tucker eds., *Contributions to The Theory of Games*, *Annals of Mathematical Studies* 28, Princeton University Press, 193-216.
5. Lazear, E. P. and Rosen, S., (1981). "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy* **89**, 841-64.
6. Loeve, (1960). *Probability Theory*. New York, Van-Nostrand.
7. Nalebuff, B. J., and Stiglitz, J. E. (1983). "Prizes and Incentives: Towards a General Theory of Compensation and Competition," *Bell Journal of Economics* **14**, 21-43.
8. Rósen, B. (1965). "Limit Theorems for Sampling," *Arkiv för Matematik* **5**, 384-424.
9. Rosen, S. (1986). "Prizes and Incentives in Elimination Tournaments," *American Economic Review* **76**, 701-15.