

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1255

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References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

COMPETITIVE PRIZES: WHEN LESS SCRUTINY INDUCES MORE EFFORT

Pradeep Dubey and Chien-wei Wu

May 2000

Competitive prizes: when less scrutiny induces more effort

May 8, 2000

Pradeep Dubey & Chien-wei Wu ¹

Abstract

We consider a principal who is keen to induce his agents to work at their maximal effort levels. To this end, he samples n days at random out of the T days on which they work, and awards a prize of B dollars to the most productive agent. The principal's policy (B, n) induces a strategic game $\Gamma(B, n)$ between the agents. We show that to implement maximal effort levels weakly (or, strongly) as a strategic equilibrium (or, as dominant strategies) in $\Gamma(B, n)$, at the least cost B to himself, the principal must choose a small sample size n . Thus less scrutiny by the principal induces more effort from the agents.

The need for reduced scrutiny becomes more pronounced when agents have information of the history of past plays in the game. There is an inverse relation between information and optimal sample size. As agents acquire more information (about each other), the principal—so to speak—must “undo” this by reducing his information (about them) and choosing the sample size n even smaller.

Journal of Economic Literature Classification Numbers: C720, D820, J410.

Key words: competitive prizes, extensive form games, information patterns, strategic equilibria, optimal sample sizes.

¹Center for Game Theory, Department of Economics, SUNY at Stony Brook

We are grateful to the Cowles Foundation for Research in Economics, Yale University, and to the Centre for Economic Studies and Planning, Jawaharlal Nehru University, where part of this work was carried out. In particular, we wish to thank Sugato Dasgupta, Ori Haimanko and Kunal Sengupta for thoughtful comments.

1 Introduction

The intuition behind this paper is simple. Suppose a principal writes an anonymous contract between several agents, basing his reward on their relative performance. To induce them to work (at their maximal effort levels), he would do well to ensure that the *least skilled* agent has a sufficiently good chance of obtaining the reward. For then this agent works hard, which generates the competition spurring the next-best agent to also work, which in turn spurs the agent just above him and so on. By designing a contract which is favorable to the weakest agent, the principal can trigger competition which elicits maximal effort from all of them. We believe that this basic principle is of quite wide application², even though for concreteness we consider here a stylized model from the principal-agent literature. In this model, it turns out that to favor the weak is tantamount to observing agents' outputs on a *small* sample, since all large samples would almost surely distinguish the weak from the strong and destroy incentives for the weak. In short, less scrutiny by the principal induces more effort from the agents.

The principal-agent literature is well-known and we make no attempt to summarize it here (see, e.g., Green & Stokey (1983), Grossman & Hart (1983), Lazer & Rosen (1981), Mookherjee (1984), Nalebuff & Stiglitz (1983) and Rosen (1986) and the references therein). Our focus is on a special scenario. There is a principal who values his agents' outputs so highly (e.g., because they fetch a very remunerative market price) that even after compensating them for any additional effort, he makes a profit on the margin. An optimal policy for the principal thus necessarily entails *maximal effort* from the agents.

Agents, on the other hand, have to be induced to work by offers of suitably large performance-related rewards of money, since they have a natural disutility for work. Their outputs are random but positively correlated with their effort levels. If an agent produces small output, it could be because he did not work much or because he had bad luck despite hard work. Of course, the probability of small output is reduced if his effort level goes up.

We suppose that the principal is constrained to write *non-discriminatory* contracts which are based on *output alone*. This could be because he cannot observe the inputs of effort made by the agents, nor can he tell them apart

²e.g., the handicap offered to the weaker opponent in golf or chess.

in terms of their private characteristics, such as their skill (probability of being productive) or their utility for money and leisure. Alternatively, even if the principal were fully cognizant of these, it could be that the law requires rewards to be commensurate with performance, and to be based solely on the outputs agents produce, with no other form of discrimination permitted. In any case, we confine attention to contracts that are defined anonymously on the outputs of the agents.

Contracts may be *independent* in that the reward to any agent depends just on his own output, i.e., it is invariant of others' outputs; or they may be *competitive* rewarding agents according to the rank-order of their outputs.

The focus of this paper is on non-discriminatory competitive contracts, which take the form of a prize to the most productive agent. Such contracts can often be more favorable to the principal than independent contracts (see section 12). But, apart from their theoretical *raison d'être*, competitive prizes are simply a fact of everyday life. Think, for example, of a bonus to the best salesman of the year; or a promotion awarded to the branch manager who produced the highest profits, etc. The question arises: how is best performance to be measured? We point out that, even if the entire stream of agents' outputs is susceptible of costless observation, there are many situations in which the principal would be wise to deliberately create uncertainty by observing outputs only on random samples of small size.

Let us consider a concrete instance of our model. Suppose a government (the principal) has commissioned two manufacturers (the agents) to produce T units each of some vital defence equipment³, e.g., timed fuses for bombs. Each manufacturer $\alpha \in \{1, 2\}$ must choose an effort level e_α from $E_\alpha \equiv \{1, \dots, e_\alpha^*\}$, which in turn determines the probability $p_\alpha^{e_\alpha}$ that any one of his fuses will be of high quality. The disutility incurred by α on account of effort e_α is $d_\alpha(e_\alpha)$. We may think of the effort as being used to upgrade the production technology. It is then natural to assume $p_\alpha^{e_\alpha} > p_\alpha^{\tilde{e}_\alpha}$ and $d_\alpha(e_\alpha) > d_\alpha(\tilde{e}_\alpha)$ whenever $e_\alpha > \tilde{e}_\alpha$.

The government buys the T fuses from 1,2 for some previously contracted amounts. But it is evidently crucial to the government that both manufacturers commit themselves to maximal effort. To motivate them, it announces an award to be bestowed as follows. It will sample n fuses at random at both production sites, and give a bonus of B dollars to the manufacturer with more high quality pieces in the sample. In the event of a tie,

³More generally, the government could commission T_1, T_2 units and sample n_1, n_2 units and award the bonus to the manufacturer with the higher percentage of high quality pieces. Our main point remains intact. At least one of n_1, n_2 must be small to avoid the need for huge bonuses. We take $T_1 = T_2 = T$ and $n_1 = n_2 = n$ for simplicity.

it will award B to each with probability $1/2$.

A policy (B, n) induces a strategic game $\Gamma(B, n)$ between the two manufacturers. The problem facing the government is how to choose (B, n) so as to minimize B while, at the same time, ensuring that (e_1^*, e_2^*) is a strategic equilibrium of $\Gamma(B, n)$. For any fixed sample size n it is clear that, if B is large enough, then (e_1^*, e_2^*) is a strategic equilibrium of $\Gamma(B, n)$. We prove somewhat more: there exists a threshold bonus $B(n)$ such that, if $B > B(n)$, then (e_1^*, e_2^*) is the *unique* strategic equilibrium (SE) of $\Gamma(B, n)$; and, if $B < B(n)$, then (e_1^*, e_2^*) is *not* an SE of $\Gamma(B, n)$.

Thus the competitive contract (B, n) implements (e_1^*, e_2^*) in a strong sense: not merely as an SE of $\Gamma(B, n)$ as in much of the literature (see, e.g., Mookherjee (1984) for a synopsis), but as its unique SE. Uniqueness overcomes the vexing uncertainty of how agents will coordinate upon the “chosen” candidate SE when there are many of them. The only uncertainty that remains is whether the agents will somehow settle down to an SE of the game at all. There is, in this context, the much stronger notion of implementation of (e_1^*, e_2^*) as a *dominant strategy* equilibrium (DSE), whereby this uncertainty is considerably reduced. But the reduction comes with a price. The threshold $\tilde{B}(n)$, needed to ensure that (e_1^*, e_2^*) is a DSE of $\Gamma(B, n)$, is generally much higher than $B(n)$.

An *optimal competitive prize* is obtained by choosing a sample size for which the threshold bonus is minimized, i.e., $(B(k), k)$ is optimal if

$$B(k) = \min \{B(n) : 1 \leq n \leq T\}.$$

It is easy to see that k must be small. Let manufacturer 1 be more skilled than 2, i.e., $p_1^{e_1^*} > p_2^{e_2^*}$. If k were big then, by the law of large numbers, 1 would almost surely produce $p_1^{e_1^*}k$ high quality fuses in any sample at (e_1^*, e_2^*) , while his rival would almost surely produce only $p_2^{e_2^*}k < p_1^{e_1^*}k$ such fuses. To induce 2 to choose e_2^* , B must be huge, compensating 2 for his extremely low probability of winning B . It follows that an optimal sample size k must be small, justifying the title of our paper. (This argument incidentally also shows that an optimal sample size \tilde{k} for implementing (e_1^*, e_2^*) as a DSE, i.e., $\tilde{B}(\tilde{k}) = \min \{\tilde{B}(n) : 1 \leq n \leq T\}$, is also small.)

The need for reduced scrutiny comes even more to the fore when agents have information of the history of past plays in the game. It is useful here to distinguish two scenarios. In the first, agents are ignorant of their rivals’ outputs⁴, but may have memory of their own outputs. In the second, they

⁴We assume throughout that no agent can observe the inputs of effort made by his rival.

can observe rivals' outputs as well. Let σ^* denote the strategies wherein all agents exert maximal effort in all circumstances of the game. We show that while all sample sizes are feasible in the first scenario (in that they implement σ^* for suitably large bonuses), *optimal* sample sizes are small and tend to fall as memory is refined. In the second scenario, even *feasible* sample sizes must be small. More precisely, there is an upper bound on feasible sample sizes, which depends on the information pattern in the game. As information is refined, this bound falls sharply. (We note, in passing, that the automatic privilege enjoyed by σ^* of being the unique SE is lost in the second scenario. The bonus needs to be raised further to bring back uniqueness.)

Two authors (Cowen & Glazer (1996)) have commented on a similar phenomenon. Their analysis has been in the context of a contract between a principal and a *single* agent. We report on their work in some detail in section 13 and contrast it with our own. But to the best of our knowledge no one has considered the impact of the principal's observation on agents' behaviour in a competitive framework.

The scope of our analysis has a natural limitation. We *fix the behaviour* of agents at σ^* (i.e., maximal effort) and examine *variable prizes* which implement σ^* . This is consistent with profit-maximization for the principal if he values agents' outputs sufficiently *and if agents' information in the game is not too fine*. Our analysis reveals that, with very fine information, it is not possible to implement σ^* without handing out huge bonuses (for all but very small sample sizes). Then it becomes more natural to take a complementary viewpoint: to *fix the prize* and examine the *variable behaviour* of agents that is induced by the prize. We do so in a companion paper (Dubey & Haimanko (2000)). An optimal policy no longer induces maximal effort in all circumstances of the game, but it still requires sample size to be small, and our main theme remains intact.

For ease of exposition, we carry out the analysis for the case of two agents. But all our results, except for Theorem 3 (on uniqueness of SE) hold for any number of agents with obvious modifications in the definitions and the proofs.

2 The Extensive Form Game (Ψ, B, n)

We present an extensive form game between the agents which involves T time periods. To keep notation simple, we assume that all the data of the game, except for agents' information, is stationary.

Each agent $\alpha \in \{1, 2\}$ has a finite set E_α of effort levels available every period, with a distinguished element e_α^* which represents his maximal effort level. Any $e \in E_\alpha$ induces a probability distribution p_α^e with *full support* on a finite set of nonnegative outputs $Q_\alpha \equiv (q_\alpha^1, \dots, q_\alpha^{m(\alpha)})$ producible by α , where we have arranged $q_\alpha^1 < \dots < q_\alpha^{m(\alpha)}$. Denote the rival of α by β , and denote $\bar{q}_\alpha \equiv q_\alpha^{m(\alpha)}$, $\underline{q}_\alpha \equiv q_\alpha^1$. We assume, for any agent α , that $\bar{q}_\alpha > \underline{q}_\beta$.

The game is played iteratively follows. In each period both agents make a simultaneous choice of effort levels which leads to random outputs, in accordance with the specified probabilities, and brings the game into period $t + 1$. Agents then choose effort levels again.

For the formal description, let $\Omega(t)$ denote the set of agents' nodes in period $t \in T \equiv \{1, \dots, T\}$. The set $\Omega(1) \equiv \{\omega^*\}$ is a singleton, and ω^* signifies the start of the game.

At any $\omega \in \Omega(t)$, both agents 1, 2 simultaneously choose effort levels from E_1, E_2 . After $(e_1, e_2) \in E_1 \times E_2$, there is a move of chance. Chance selects $(q_1, q_2) \in Q_1 \times Q_2$ with positive probability⁵ $p^{e_1}(q_1)p^{e_2}(q_2)$, leading to an agents' node $\bar{\omega} \equiv (\omega, e_1, e_2, q_1, q_2)$ in $\Omega(t + 1)$. See Figure 1.




Figure 1

Thus, for $t > 1$, $\Omega(t)$ is isomorphic to $(E_1 \times E_2 \times Q_1 \times Q_2)^{t-1}$; and any $\omega \in \Omega(t)$ is specified by its history $(e_1(\tau), e_2(\tau), q_1(\tau), q_2(\tau))_{\tau=1}^{t-1}$.

Let $\Omega \equiv \coprod_{t \in T} \Omega(t)$ be the set of all agents' nodes in the game tree (where \coprod denotes disjoint union). Each agent α has an *information partition* I_α of Ω which reflects what he knows of his own and his rival's past history in the game. We assume throughout that *any agent α can observe only the outputs produced by his rival and not the rival's inputs of effort*, i.e., for any pair of nodes $\omega, \bar{\omega}$ in $\Omega(t) \subset \Omega$ and any agent α in $\{1, 2\}$:

⁵i.e., agents' probabilities are independent (for correlated probabilities, see Section 11).

$$(I) \left\{ \begin{array}{l} \omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \\ \tilde{\omega} \equiv (e_\alpha(\tau), \tilde{e}_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \\ \text{i.e., } \omega \text{ and } \tilde{\omega} \text{ have the same history} \\ \text{except perhaps for the effort of } \beta \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} \omega \text{ and } \tilde{\omega} \text{ are in} \\ \text{the same information} \\ \text{set of } I_\alpha \end{array} \right\}$$

There is a positive correlation between effort and productivity. Our key assumption on p^{e_α} states that when an agent exerts maximal effort, his outputs go up, in the sense of first-order stochastic dominance:

For any $\alpha \in \{1, 2\}$ and $e \in E_\alpha \setminus \{e_\alpha^*\}$,

$$(II) \quad \sum_{q \geq x} p_\alpha^{e_\alpha^*}(q) > \sum_{q \geq x} p_\alpha^e(q) \text{ for all } x \in Q_\alpha \setminus \{q_\alpha\}.$$

This completes the description of the extensive form of the game, but it still remains to specify agents' payoffs at each terminal node $\omega \in \Omega(T+1)$.

Each agent α has a continuous, strictly monotonic utility for money given by $u_\alpha : \mathbf{R}_+ \rightarrow \mathbf{R}$, where $u_\alpha(B) \equiv$ the utility to α of receiving B dollars by way of reward the end of the T days. We need to assume that agents value money sufficiently. Precisely (and, for the sake of simplicity, stating it in a stronger form than necessary) we postulate:

$$(III) \quad u_\alpha(B) \rightarrow \infty \text{ as } B \rightarrow \infty.$$

The disutility for effort is given by a function $d_\alpha : (E_\alpha)^T \rightarrow \mathbf{R}_+$, for $\alpha \in \{1, 2\}$. We assume that maximal effort incurs the most disutility, i.e.,

$$(IV) \quad d_\alpha((e_\alpha^*, \dots, e_\alpha^*)) > d_\alpha(e)$$

for all $e \in (E_\alpha)^T \setminus \{(e_\alpha^*, \dots, e_\alpha^*)\}$.

Rewards at the terminal nodes are determined by the *principal's policy* (B, n) which operates as follows: n days are sampled at random⁶ and the bonus of B dollars is awarded to the agent with the higher total output across the days in the sample; in the event of a tie, it is given to each agent with probability half. The principal samples secretly. Agents know that exactly n days are sampled, but not know *which* n ; this is revealed to them only ex post at the end of the game.

For $1 \leq n \leq T$, let $\mathcal{C}_n \equiv \{\mathfrak{S} \subset T : |\mathfrak{S}| = n\}$ and $P(n) \equiv 1/C(T, n)$, where $C(T, n) = T!/(T-n)!n!$. Consider a terminal node $\omega \equiv (e_1(\tau), e_2(\tau), q_1(\tau), q_2(\tau))_{\tau=1}^T \in$

⁶We will assume that the *same* set of n days is sampled for each agent. But, given our stationary set up, not much changes if these sets are sampled independently for the two agents.

$\Omega(T + 1)$. Given ω , the total produced by α on $\mathfrak{S} \in \mathcal{C}_n$, is $Z_\alpha(\omega, \mathfrak{S}) \equiv \sum_{\tau \in \mathfrak{S}} q_\alpha(\tau)$. Recall that β denotes the rival of $\alpha \in \{1, 2\}$, and define⁷

$$u_\alpha(\omega, \mathfrak{S}) \equiv \begin{cases} u_\alpha(B) & \text{if } Z_\alpha(\omega, \mathfrak{S}) > Z_\beta(\omega, \mathfrak{S}) \\ u_\alpha(0) = 0 & \text{if } Z_\alpha(\omega, \mathfrak{S}) < Z_\beta(\omega, \mathfrak{S}) \\ \frac{1}{2}u_\alpha(B) & \text{if } Z_\alpha(\omega, \mathfrak{S}) = Z_\beta(\omega, \mathfrak{S}) \end{cases}$$

Then, abusing the notation u_α yet again, the expected utility from money at ω is

$$u_\alpha(\omega) \equiv P(n) \sum_{\mathfrak{S} \in \mathcal{C}_n} u_\alpha(\omega, \mathfrak{S})$$

On the other hand, the disutility of effort at ω is

$$d_\alpha(\omega) \equiv d_\alpha((e_\alpha(\tau))_{\tau=1}^T)$$

We assume separability, and write the payoff to agent α at ω as

$$\pi_\alpha(\omega) \equiv u_\alpha(\omega) - d_\alpha(\omega)$$

This completes the description of the extensive form game between the agents. Since we hold u_α, d_α fixed throughout, we will denote the game (Ψ, B, n) where Ψ is the extensive form of the game, i.e., the game tree with the payoffs missing at the terminal nodes.

Remark 1 The example of “fuses” given in the introduction, can be seen to be a special case of our model by setting $I_1 = I_2 = \Omega$. (There is no loss of generality in ignoring the fixed contracted payment to the manufacturers, since affine transformations of their utilities leaves only the bonus term.)

When $I_1 = I_2 = \Omega$, we say that there is *zero-information* in the game. In this case, our theorems hold without the full-support assumption: $p_\alpha^e(q) > 0$ for all $q \in Q_\alpha$ and all $e \in E_\alpha$.

Remark 2 When $I_1 = I_2 = \{\Omega(t) : 1 \leq t \leq T\}$, we say that there is *low-information* in the game. Agents now know only the time period they are in. Low information may be interpreted in terms of a *locations model*. Suppose the principal hires agents 1, 2 to work for him in two isomorphic sets of locations $1, \dots, T$ and $1', \dots, T'$ respectively, and that all the work has to be carried out simultaneously in the different locations. (One could think of a government that has appointed two service-providers, and wants

⁷W.l.o.g. (using affine transformations for each agent's utility), we take $u_\alpha(0) = 0$, to avoid writing $(1/2)[u_\alpha(0) + u_\alpha(B)]$ in place of $(1/2)u_\alpha(B)$

them to undertake one-shot improvements in their respective territories; or of a manufacturer who has engaged two distributors to operate in disjoint locations.) The principal samples n corresponding locations from each set and awards B dollars to the agent who is more productive on the sample.

3 The Normal Form Game $\Gamma(\Psi, B, n)$

We next describe the game in normal (or strategic) form that arises from the extensive form game.

A strategy σ_α of agent α specifies an effort level $\sigma_\alpha(\omega)$ in E_α for each $\omega \in \Omega$, and is measurable with respect to α 's information partition I_α .

It will be useful to define an equivalence relation on the set of all strategies of agent α , so that any two strategies which have the same “reduced form” get identified. To this end, we first define the set of irrelevant nodes in $\Omega(t)$ for strategy σ_α , inductively on t . At $t = 1$, the set of irrelevant nodes is the empty set. Let $\tilde{\omega} = (\omega, e_1, e_2, q_1, q_2)$ be a node in $\Omega(t + 1)$ that follows from ω in $\Omega(t)$. Then define $\tilde{\omega}$ to be irrelevant for σ_α if either $\omega \in \Omega(t)$ is irrelevant for σ_α ; or if $\sigma_\alpha(\omega) \neq e_\alpha$. Denote the set of all irrelevant nodes of σ_α by $\Omega^{\sigma_\alpha} \subset \Omega$. We say that $\sigma_\alpha \approx \tilde{\sigma}_\alpha$ iff $\Omega^{\sigma_\alpha} = \Omega^{\tilde{\sigma}_\alpha}$ and $\sigma_\alpha(\omega) = \tilde{\sigma}_\alpha(\omega)$ for all $\omega \in \Omega \setminus \Omega^{\sigma_\alpha}$.

It is clear that \approx is an equivalence relation, and *from now on a strategy will be thought of as an equivalence class*.

The set of all (reduced as above) strategies of agent α is denoted Σ_α .

Any pair of *unreduced* strategies $\sigma \equiv (\sigma_1, \sigma_2)$ induces a probability distribution on the set of terminal nodes in the obvious manner. Indeed, let $\tilde{\omega} \equiv (\omega, e_1, e_2, q_1, q_2)$ be an immediate follower of ω in the game tree, and define the probability $p^\sigma(\omega, \tilde{\omega})$ of reaching $\tilde{\omega}$ from ω under σ to be $p_1^{e_1}(q_1)p_2^{e_2}(q_2)$ if $\sigma_\alpha(\omega) = e_\alpha$ for both $\alpha = 1$ and $\alpha = 2$, and to be 0 otherwise. Then the probability⁸ $p^\sigma(\omega)$ of reaching the node $\omega = (e_1(\tau), e_2(\tau), q_1(\tau), q_2(\tau))_{\tau=1}^{\tilde{t}}$ is $p^\sigma(\omega) = \prod_{t=1}^{\tilde{t}} p^\sigma(\omega_t, \omega_{t+1})$ where $\omega_{t+1} \equiv (\omega_t, e_1(t), e_2(t), q_1(t), q_2(t))$ for $t = 1, \dots, \tilde{t}$. (We define $p^\sigma(\omega^*) = 1$ where, recall, ω^* is the start of the game.) The payoff to agent α from σ is then $\sum_{\omega \in \Omega(T+1)} p^\sigma(\omega) \pi_\alpha(\omega) \equiv \Pi_\alpha(\sigma)$.

If (σ_1, σ_2) and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ are two equivalent pairs (i.e., $\sigma_\alpha \approx \tilde{\sigma}_\alpha$ for $\alpha \in \{1, 2\}$), it is easy to check that they induce the same probability distribution on $\Omega(T + 1)$. Thus equivalent pairs give rise to the same expected payoffs

⁸Note: $p^{(\sigma_1, \sigma_2)}(\omega)$ is in fact defined for arbitrary maps $\sigma_\alpha : \Omega \rightarrow E_\alpha$, *not* necessarily just strategies $\sigma_\alpha \in \Sigma_\alpha$

and the normal form game, with strategy-sets Σ_1 and Σ_2 , is well-defined. We denote it by $\Gamma(\Psi, B, n)$.

Let $\sigma_\alpha \in \Sigma_\alpha$ and $\sigma_\beta \in \Sigma_\beta$. Then σ_α is called a *best reply* to σ_β if

$$\Pi_\alpha(\sigma_\alpha, \sigma_\beta) \geq \Pi_\alpha(\tilde{\sigma}_\alpha, \sigma_\beta) \text{ for all } \tilde{\sigma}_\alpha \in \Sigma_\alpha$$

The pair (σ_1, σ_2) is called a *strategic equilibrium* (and denoted SE) of $\Gamma(\Psi, B, n)$ if σ_1 is a best reply to σ_2 , and σ_2 is a best reply to σ_1 . If, moreover, both strategies are unique best replies then the SE is called *strict*.

Finally we define $\sigma_\alpha \in \Sigma_\alpha$ to be a *dominant (strictly dominant) strategy* of α in $\Gamma(\Psi, B, n)$ if σ_α is a best reply (unique best reply) to *every* $\sigma_\beta \in \Sigma_\beta$. If σ_1 and σ_2 are both dominant strategies in $\Gamma(\Psi, B, n)$, we say that (σ_1, σ_2) is a dominant strategy equilibrium (DSE) of $\Gamma(\Psi, B, n)$.

4 Thresholds for the Bonus

The strategy (modulo the equivalence relation) wherein α *puts in maximal effort* is denoted σ_α^* , i.e.,

$$\sigma_\alpha^*(\omega) = e_\alpha^* \text{ for all } \omega \in \Omega.$$

We denote $(\sigma_1^*, \sigma_2^*) \equiv \sigma^*$.

First let us fix the extensive form Ψ and the sample size n , and ask how large the bonus needs to be to implement σ^* either as an SE or as a DSE. Define

$$\begin{aligned} B(\Psi, n) &\equiv \min\{B : \sigma^* \text{ is an SE of } \Gamma(\Psi, B, n)\} \\ \tilde{B}(\Psi, n) &\equiv \min\{B : \sigma^* \text{ is a DSE of } \Gamma(\Psi, B, n)\} \end{aligned}$$

(The minimum on an empty set is taken to be infinity.)

It is straightforward to check

Theorem 1

- (i) $B > B(\Psi, n) \Leftrightarrow (\sigma_1^*, \sigma_2^*)$ is a strict SE of $\Gamma(\Psi, B, n)$
- (ii) $B > \tilde{B}(\Psi, n) \Leftrightarrow \sigma_\alpha^*$ is a strictly dominant strategy of $\Gamma(\Psi, B, n)$ for $\alpha \in \{1, 2\}$

This theorem justifies the use of the word “threshold” for $B(\Psi, n)$ and $\tilde{B}(\Psi, n)$.

5 Feasible and Optimal Sample Sizes

Since the principal is interested in implementing σ^* with the smallest possible bonus, it is useful to make some definitions.

A sample size n is called

- (i) *SE-feasible* for Ψ if $B(\Psi, n) < \infty$
- (ii) *DSE-feasible* for Ψ if $\tilde{B}(\Psi, n) < \infty$
- (iii) *SE-optimal* for Ψ if $B(\Psi, n) = \min\{B(\Psi, k) : 1 \leq k \leq T\} < \infty$
- (iv) *DSE-optimal* for Ψ if $\tilde{B}(\Psi, n) = \min\{\tilde{B}(\Psi, k) : 1 \leq k \leq T\} < \infty$

6 Ignorance of the Rival: Memory and Sample Size

We shall say that agent α is *ignorant of his rival* β in Ψ if, for any pair of nodes $\omega, \tilde{\omega}$ in $\Omega(t) \subset \Omega$ and any agent α ,

$$\left\{ \begin{array}{l} \omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \\ \tilde{\omega} \equiv (e_\alpha(\tau), \tilde{e}_\beta(\tau), q_\alpha(\tau), \tilde{q}_\beta(\tau))_{\tau=1}^{t-1} \\ \text{i.e., } \omega \text{ and } \tilde{\omega} \text{ have the same history} \\ \text{of the efforts and outputs of agent } \alpha \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} \omega \text{ and } \tilde{\omega} \text{ are in} \\ \text{the same information} \\ \text{set of } I_\alpha \end{array} \right\}$$

When each agent is ignorant of his rival, σ^* can be sustained as an SE for many sample sizes with finite (albeit large) bonus. Indeed define $\tilde{\Delta} \equiv \max\{|\underline{q}_1 - \underline{q}_2|, |\bar{q}_1 - \bar{q}_2|\}$, $\Delta \equiv \min\{\bar{q}_1 - \underline{q}_2, \bar{q}_2 - \underline{q}_1\}$. Then we have

Theorem 2 (Existence of SE)

Suppose each agent is ignorant of his rival in Ψ . Then, if $n \leq (\Delta/\tilde{\Delta}) + 1$, n is SE-feasible and DSE-feasible for Ψ .

Corollary Suppose each agent is ignorant of his rival in Ψ and $\bar{q}_1 = \bar{q}_2$ and $\underline{q}_1 = \underline{q}_2$. Then all sample sizes are SE-feasible and DSE-feasible for Ψ . (This follows since $\tilde{\Delta} = 0$.)

Whenever the bonus is higher than the threshold, σ^* is the *unique* SE.

Theorem 3 (Uniqueness of SE)

Suppose each agent is ignorant of his rival in Ψ . Let n be any feasible sample size for Ψ . Then

$$B > B(\Psi, n) \Rightarrow \sigma^* \text{ is the unique SE of } \Gamma(\Psi, B, n)$$

(In fact, σ^* is the unique SE in the mixed extension of $\Gamma(\Psi, B, n)$.)

According to Theorem 3, whenever the principal can implement σ^* as an SE at all, he need not further worry about the vexing issue of multiple equilibria: adding a penny to the threshold will make σ^* the unique SE. (It is only at the threshold $B(\Psi, n)$ that σ^* may have to coexist with other SE of the game.)

Suppose memory is refined in the game, i.e., I_α is refined to \tilde{I}_α , still leaving α ignorant of his rival. This will tend to reduce optimal sample sizes. The intuition is clear. Suppose an agent has succeeded on most days over a long past and knows so on account of refined memory. If the sample size is large, he will tend to shirk on the remaining few days because this hardly affects his probability of winning the bonus. To get him to work there, B would have to be huge. But if the sample size is small, he gives up a good chance of winning B by shirking. Thus a smaller B will sustain full effort once the sample size is lowered. Indeed, in the extreme case when the sample size is 1, the increase in the probability of winning through an extra day's work remains invariant of memory.

It might help to see a concrete calculation.

Example 1

Suppose each agent α has just two effort levels and two output levels denoted $E_\alpha \equiv \{0, 1\}$, $Q_\alpha \equiv \{0, 1\}$. Let $p_1^1(1) = 0.7$, $p_1^0(1) = \epsilon$ and $p_2^1(1) = 0.6$, $p_2^0(1) = \epsilon$ for small ϵ . Thus agent 1 is more "skilled" than his rival, in that he produces output 1 with higher probability when he works. Suppose there are 180 periods, i.e., $T = \{1, \dots, 180\}$ and low information, i.e., $I_1 = I_2 = \{\Omega(t) : 1 \leq t \leq 180\}$. For any sequence of efforts $e \in \{0, 1\}^{180}$, let $|e| \equiv$ number of 1's in the sequence, i.e., the number of days of maximal effort. Put $d_1(e) = 2|e|$, $d_2(e) = |e|$. Finally let $u_1(x) = u_2(x) = x$. This defines the extensive form game between the two agents. The optimal sample size is 26, as we show in Section 8.

Now suppose both agents have perfect memory (while remaining ignorant of each other). Then the optimal sample size must be 2. To see this, consider any sample size $1 \leq k \leq 180$ and suppose the last day has arrived, with the

weak agent 2 fully aware that he has produced 0 on all the previous 179 days. This is the worst scenario in which to incentivize him to work. (If the bonus makes him work here, it will make him work at all other positions in the game tree.) It is clear that he will not work if he lags by more than 1 unit (in total output) over the previous days of the sample. So define, for $j = 0$ or 1 , $R_j(k) =$ probability that the rival agent 1 produces exactly j units across the previous random $k - 1$ days of the sample (given that 1 is putting in maximal effort on all days). Let $P(k)$ be the probability that the last day is in the sample. Then, by switching from effort 0 to 1 on the last day, agent 2 increases the probability of winning the bonus by

$$\Delta(k) = P(k)[R_1(k)f_1(p_1, p_2) + R_0(k)f_2(p_1, p_2)]$$

where $f_1(p_1, p_2) = [(1/2)(1 - p_1^1)(p_2^1 - p_2^0)]$ and $f_2(p_1, p_2) = (1 - p_1^1)[(p_2^1 - p_2^0) + (1/2)((1 - p_2^1) - (1 - p_2^0))] + p_1^1[(1/2)(p_2^1 - p_2^0)]$

The bonus B needed to make 2 work on the last day must satisfy

$$\Delta(k)u_2(B) \geq 1 \equiv 2\text{'s disutility to work on the last day.}$$

Now $P(k) = k/T$ increases linearly with k , while both $R_j(k)$ fall geometrically. Thus for all large enough k , $R_j(k)$ will be very small necessitating a huge bonus, much bigger than the bonus corresponding to sample size 1. In our example, it is easy to calculate that this happens for $k \geq 3$. On the other hand, for small k , the rise of $P(k)$ may dominate the fall of $R_j(k)$, so that $\Delta(k)$ increases with k . In our example this occurs when we go from $k = 1$ to $k = 2$, explaining why sample size 2 is always better than sample size 1, and is the optimal sample size. (Indeed this is so for all $T \geq 3$.)

7 Information of the Rival and Sample Size

When agents have information regarding each other's outputs in Ψ , a dramatic change occurs. Feasible sample sizes have an upper bound (i.e., for sample sizes that exceed the bound, no amount of bonus can sustain σ^* as an SE). There is an intimate inverse relationship between the information pattern in the game tree and this bound. As information is refined, the bound falls. To put it another way: in order to restore σ^* as an SE, the principal must "undo" the gain in agents' information (about each other) by reducing his own information (about them) via a smaller sample size.

We shall assume throughout this section that an agent has perfect recall. Furthermore his information of past outputs, achieved by him and his rival,

is invariant of his memory of his own past efforts⁹. (This is tenable since random outputs have the same support for every effort level.) To be precise, let us define $\omega \sim (I_\alpha)\tilde{\omega}$ if ω and $\tilde{\omega}$ are in the same information set in I_α . Also for $\omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1}$ and $1 \leq t' \leq t-1$, denote $(\omega|t') \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t'}$. Keeping in mind that (see (I)) that α does not observe his rival's efforts, we postulate:

(V) (i) **(Perfect recall)**

$$\left\{ \begin{array}{l} \omega \sim (I_\alpha)\tilde{\omega}; \\ \omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1}; \\ \tilde{\omega} \equiv (\tilde{e}_\alpha(\tau), \tilde{e}_\beta(\tau), \tilde{q}_\alpha(\tau), \tilde{q}_\beta(\tau))_{\tau=1}^{t-1}; \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} t = \tilde{t}; \\ e_\alpha(\tau) = \tilde{e}_\alpha(\tau) \text{ for } 1 \leq \tau \leq t-1; \\ (\omega|t') \sim (I_\alpha)(\tilde{\omega}|t') \text{ for } 1 \leq t' < t-1. \end{array} \right\}$$

(ii) **(Invariance of Output Information on Memory of Effort)**

There exist partitions $J_\alpha(\tilde{t})$ of $Q_\alpha^{\tilde{t}-1} \times Q_\beta^{\tilde{t}-1}$ for $1 < \tilde{t} \leq T$ which characterize I_α in the following sense: $S \in I_\alpha$ implies

$$\exists t \in T, (e_\alpha(\tau))_{\tau=1}^{t-1} \in E_\alpha^{t-1}, K \in J_\alpha(t) \text{ such that}$$

$$S = \{(e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} : e_\beta(\tau) \in E_\beta, (q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \in K\}$$

We shall need one more assumption for this section.

(VI) **(Sufficiently Fine Grid on Outputs)** For any $\alpha \in \{1, 2\}$, $q_\beta^{l+1} - q_\beta^l < \bar{q}_\alpha - \underline{q}_\alpha$ for $1 \leq l < m(\beta)$.

This is not too restrictive, e.g., it automatically holds if the sets $Q_\alpha = \{\underline{q}_\alpha \equiv q_\alpha^1, \dots, q_\alpha^{m(\alpha)} \equiv \bar{q}_\alpha\}$ consist of consecutive integers, for $\alpha \in \{1, 2\}$.

We define a positive integer $f_\alpha(\omega)$ at each node $\omega \in \Omega$ and for each agent α . The functions f_α do not depend on the information pattern or the probabilities of production in Ψ . Intuitively, $f_\alpha(\omega)$ is the largest sample size which still leaves some incentive for agent α to put in maximal effort e_α^* at the node ω , *when he knows ω perfectly*.

Then we define, for the extensive form Ψ ,

$$g(\Psi) = \min_{\alpha \in \{1, 2\}} \min_{S \in I_\alpha} \max \{f_\alpha(\omega) : \omega \in S\}$$

⁹There are many situations in which agents may possess perfect information of the history of their own past efforts, but not of the outputs they have produced. Suppose, for instance, that there is time lag in production, and that the output corresponding to day t 's effort comes to light only some days later.

Thus g assigns a positive integer to every extensive form Ψ . In contrast to the functions f_α , the function g is highly sensitive to the information in Ψ .

Pending the precise definition of f_α to the proof of Theorem 4 in Section 17.2.1, we are ready to state

Theorem 4 Assume¹⁰ I, II, III, IV, V, VI. Then

$$\text{the sample size } n \text{ is SE-feasible for } \Psi \Leftrightarrow n \leq g(\Psi).$$

Theorem 4 makes precise our intuition that, in order to sustain σ^* as an SE with finite bonus, the principal must lower his sample size when agents' information is increased. Let us say that $\tilde{\Psi}$ is an information-refinement of Ψ (and write $\tilde{\Psi} \succ \Psi$) if, for each agent α , his information partition \tilde{I}_α in $\tilde{\Psi}$ is a refinement of his information partition I_α in Ψ ; and, other than that, the extensive forms Ψ and $\tilde{\Psi}$ are identical. It is evident from our formula for g that

$$\tilde{\Psi} \succ \Psi \Rightarrow g(\tilde{\Psi}) \leq g(\Psi)$$

In general $g(\tilde{\Psi})$ falls sharply below $g(\Psi)$ when the refinement is significant. This says that if agents have more information about each other, then the principal must reduce his sample size in order to elicit maximal effort from them at an SE. Indeed, it is easy to check that if the game has three or more periods and if information is at its finest in Ψ (i.e., agents observe everything except each other's efforts), then $g(\Psi) = 2$. The need for reduced scrutiny is then indeed obvious!

Remark 3 When agents can observe each others' outputs, Theorem 3 is no longer true. We nevertheless can sustain σ^* as a unique SE for sufficiently large bonus. Indeed, for any sample size $n \leq g(\Psi)$, define

$$B_*(\Psi, n) \equiv \inf \{B \in R_+ : \sigma^* \text{ is the unique SE of } \Gamma(\Psi, B, n) \text{ for all } \tilde{B} \geq B\}$$

Then $B_*(\Psi, n) < \infty$ as we prove in Section 17. (Of course, typically $B_*(\Psi, n) \gg B(\Psi, n)$.)

¹⁰Assumptions V and VI are invoked *only* in Theorem 4. All the other results invoke only our general assumptions I, II, III, IV; except for Theorem 5 to follow, which needs Assumption VII and VIII in addition to I, II, III, IV.

8 The Principal's Optimal Policy

Define, for $\alpha \in \{1, 2\}$,

$$B_\alpha(\Psi, n) = \min \{B : \sigma_\alpha^* \text{ is a best reply to } \sigma_\beta^* \text{ in the game } \Gamma(\Psi, B, n)\},$$

Then

$$B(\Psi, n) = \max \{B_1(\Psi, n), B_2(\Psi, n)\},$$

An optimal SE-policy $(B(\Psi, k), k)$ in Ψ must satisfy

$$B(\Psi, k) = \min \{B(\Psi, n) : 1 \leq n \leq T\}.$$

(In the same vein, an optimal DSE-policy $(\tilde{B}(\Psi, \tilde{k}), \tilde{k})$ in Ψ must satisfy $\tilde{B}(\Psi, \tilde{k}) = \min \{\tilde{B}(\Psi, n) : 1 \leq n \leq T\}$ where $\tilde{B}(\Psi, n) = \max\{\tilde{B}_1(\Psi, n), (\tilde{B}_2(\Psi, n))\}$ and $\tilde{B}_\alpha(\Psi, n) = \min\{B : \sigma_\alpha^* \text{ is a dominant strategy in } \Gamma(\Psi, B, n)\}.$)

Consider again Example 1 of Section 5 in the low information case. For small ϵ , we obtain the computer-aided graph of Figure 1. (The calculations were done for $\epsilon = 0.001$).

Notice that $B_1(n)$ is downward-sloping, while $B_2(n)$ has a U-shape. As was said before, raising n increases the probability of sampling any particular day, and therefore tends to lower the bonus at which agents will work. But as n increases, the probability that the weak agent has been strongly overtaken by his rival (on the other $n - 1$ days of the sample) is high. This tends to raise the required bonus for the weak agent and to lower it for the strong agent. The two effects on the bonus go in the same direction for the strong agent, hence $B_1(n)$ is downward-sloping. However they work in opposite directions for the weak agent and so $B_2(n)$ turns and starts rising, when the second effect dominates the first. Indeed the turn comes rather quickly at $n = 3$ in our example. (As p_2 rises, keeping p_1 fixed, the turn occurs at larger sample sizes.)




Figure 2

9 Optimal Sample Size for Long Horizons

Let $\Psi(1), \Psi(2), \dots$ be a sequence of extensive forms of length $1, 2, \dots$ with fixed underlying $E_\alpha, Q_\alpha, p_\alpha$ for $\alpha \in \{1, 2\}$. Let the information partitions $I_1(T), I_2(T)$ in $\Psi(T)$ be arbitrary for $T = 1, 2, \dots$. Denote the utility for money and the disutility for work in $\Psi(T)$ by u_α^T, d_α^T . We assume that the disutility of working all the time increases linearly with T and is distributed in a non-skewed manner between the periods:

(VII) there exists $\gamma > 0$ such that, for $\alpha \in \{1, 2\}$ and any T we have

- (i) $d_\alpha^T((e_\alpha^*(\tau))_{\tau=1}^T) > T/\gamma$
- (ii) Consider $(e_\alpha^*(\tau))_{\tau=1}^T$ and $(\tilde{e}_\alpha^*(\tau))_{\tau=1}^T$. Suppose, for some $t \in T$, $e_\alpha(t) = e_\alpha^* \neq \tilde{e}_\alpha(t)$, and $e_\alpha(\tau) = \tilde{e}_\alpha(\tau)$ for all $\tau \in T \setminus \{t\}$. Then $d_\alpha^T((e_\alpha(\tau))_{\tau=1}^T) - d_\alpha^T((\tilde{e}_\alpha(\tau))_{\tau=1}^T) < \gamma$

Finally assume that agent 1 is more skilled than agent 2:

$$(VIII) \quad \sum_{q \in Q_1} p_1^{e_1^*}(q)q > \sum_{q \in Q_2} p_2^{e_2^*}(q)q$$

i.e., if both agents put in maximal effort, 1 produces a strictly higher average than 2. Define¹¹

$$B(n, T) \equiv \min \{B : \sigma^* \text{ is an SE of the game } \Gamma(\Psi(T), B, n)\}$$

and suppose

$$n(T) \in \text{Argmin}_{1 \leq n \leq T} B(n, T).$$

Then we have, justifying the title of this paper,

Theorem 5 Assume VII and VIII, and that each T -period game satisfies I, II, III, IV. Then

$$\frac{n(T)}{T} \rightarrow 0 \text{ as } T \rightarrow \infty.$$

10 Promotion or Reputation

When the reward for producing more is not a variable amount of money (i.e., a bonus of B dollars), but a fixed prize such as a promotion or enhanced reputation, we obtain a variant model. But the key to its analysis is provided by the bonus-model we have studied so far.

Consider again the special example of Section 7 with low information. Let u_α^* be the utility to α of the fixed prize, for $\alpha \in \{1, 2\}$, and by scaling utilities assume w.l.o.g. that $u_1^* = u_2^*$.

Suppose¹² $u_\alpha^* > u_\alpha(B(n))$ for $\alpha \in \{1, 2\}$ and for some n . Then, by sampling an optimal size k , the principal can ensure that both agents will put in maximal effort. Of course he can achieve the same effect by choosing any sample size n for which $u_\alpha^* > u_\alpha(B(n))$ for $\alpha \in \{1, 2\}$. But the set of such n will always contains every optimal size k . (See Figure 3, where the u_α^* have been superimposed on Example 1 and where this set forms an interval.). Thus k is optimal in the promotions-model as well.

¹¹The strategy $\sigma^*(T)$, of always putting in maximal effort in $\Psi(T)$, is denoted σ^* without confusion.

¹²Otherwise the promotion cannot induce the agents to put in maximal effort.




Figure 3

So, in the promotions-model too, the advantages to be gained from reduced scrutiny are quite evident. (For an indepth discussion of promotions, see Dubey & Haimanko (2000).)

11 Correlated Probabilities

We can permit *weak* correlation in agents' productivity without affecting many of our results. Indeed suppose that when agents choose $(e_1, e_2) \in E_1 \times E_2$ at $\omega \in \Omega(t)$, this gives rise to a probability distribution $p^{(e_1, e_2)}$ with full support on $Q_1 \times Q_2$. (Notice that while we have dropped independence of probabilities across agents at any given time, we still maintain independence across time.) The assumption that more effort boosts productivity now takes the following form. For any $\alpha \in \{1, 2\}$, $e \in E_\alpha \setminus \{e_\alpha^*\}$, $\hat{q} \in Q_\alpha \setminus \{\underline{q}_\alpha\}$ and $\tilde{q}_\beta \in Q_\beta$:

$$\sum_{q \in Q_\alpha, q \geq \hat{q}} p^{e_\alpha^*, e_\beta^*}(q, \tilde{q}_\beta) > \sum_{q \in Q_\alpha, q \geq \hat{q}} p^{e_\alpha, e_\beta^*}(q, \tilde{q}_\beta)$$

This says, in effect, that when agent α shifts to his maximal effort he boosts his own productivity, given any fixed quantity of his rival.

With this hypothesis, Theorems 1,2 and 4 still hold. Theorem 5 holds if we assume that agent α is more skilled than β in the sense:

$$\sum_{(x,y) \in Q_\alpha \times Q_\beta} p^{\epsilon_\alpha^*, \epsilon_\beta^*}(x, y)(x - y) > 0$$

(The proofs require only obvious changes in notation.) But Theorem 3 breaks down, except for $T = 1$. Notice that the basic compulsions for reducing scrutiny remain intact (by Theorems 4 and 5).

12 Optimality of Competitive Prizes

Both Green & Stokey (1983) and Mookherjee (1984) provide examples when competitive contracts are better for the principal than independent contracts. We will illustrate the same phenomenon with a simple example based on our competitive prizes. There is just one day, i.e., $T = n = 1$. Each agent $\alpha \in \{1, 2\}$ has two effort levels, corresponding to low or high effort, and denoted 0 or 1; and each can produce 0, 0.1 or 1 units of output. Probabilities of production are given by the tables

$(\sigma_1, \sigma_2) = (1, 1)$	$(\sigma_1, \sigma_2) = (1, 0)$																																
<table style="width: 100%; border-collapse: collapse;"> <tr><td style="border: none;"></td><td style="border: none; text-align: center;">0</td><td style="border: none; text-align: center;">0.1</td><td style="border: none; text-align: center;">1</td></tr> <tr><td style="border: none; text-align: center;">0</td><td style="border: 1px solid black; text-align: center;">0.8-8 ϵ</td><td style="border: 1px solid black; text-align: center;">0.5 ϵ</td><td style="border: 1px solid black; text-align: center;">1.5 ϵ</td></tr> <tr><td style="border: none; text-align: center;">0.1</td><td style="border: 1px solid black; text-align: center;">0.5 ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">1.5 ϵ</td></tr> <tr><td style="border: none; text-align: center;">1</td><td style="border: 1px solid black; text-align: center;">1.5 ϵ</td><td style="border: 1px solid black; text-align: center;">1.5 ϵ</td><td style="border: 1px solid black; text-align: center;">0.2</td></tr> </table>		0	0.1	1	0	0.8-8 ϵ	0.5 ϵ	1.5 ϵ	0.1	0.5 ϵ	ϵ	1.5 ϵ	1	1.5 ϵ	1.5 ϵ	0.2	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="border: none;"></td><td style="border: none; text-align: center;">0</td><td style="border: none; text-align: center;">0.1</td><td style="border: none; text-align: center;">1</td></tr> <tr><td style="border: none; text-align: center;">0</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">0.1</td><td style="border: 1px solid black; text-align: center;">1-8 ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">1</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> </table>		0	0.1	1	0	ϵ	ϵ	ϵ	0.1	1-8 ϵ	ϵ	ϵ	1	ϵ	ϵ	ϵ
	0	0.1	1																														
0	0.8-8 ϵ	0.5 ϵ	1.5 ϵ																														
0.1	0.5 ϵ	ϵ	1.5 ϵ																														
1	1.5 ϵ	1.5 ϵ	0.2																														
	0	0.1	1																														
0	ϵ	ϵ	ϵ																														
0.1	1-8 ϵ	ϵ	ϵ																														
1	ϵ	ϵ	ϵ																														
$(\sigma_1, \sigma_2) = (0, 0)$	$(\sigma_1, \sigma_2) = (0, 1)$																																
<table style="width: 100%; border-collapse: collapse;"> <tr><td style="border: none;"></td><td style="border: none; text-align: center;">0</td><td style="border: none; text-align: center;">0.1</td><td style="border: none; text-align: center;">1</td></tr> <tr><td style="border: none; text-align: center;">0</td><td style="border: 1px solid black; text-align: center;">1-8 ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">0.1</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">1</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> </table>		0	0.1	1	0	1-8 ϵ	ϵ	ϵ	0.1	ϵ	ϵ	ϵ	1	ϵ	ϵ	ϵ	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="border: none;"></td><td style="border: none; text-align: center;">0</td><td style="border: none; text-align: center;">0.1</td><td style="border: none; text-align: center;">1</td></tr> <tr><td style="border: none; text-align: center;">0</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">1-8ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">0.1</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> <tr><td style="border: none; text-align: center;">1</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td><td style="border: 1px solid black; text-align: center;">ϵ</td></tr> </table>		0	0.1	1	0	ϵ	1-8 ϵ	ϵ	0.1	ϵ	ϵ	ϵ	1	ϵ	ϵ	ϵ
	0	0.1	1																														
0	1-8 ϵ	ϵ	ϵ																														
0.1	ϵ	ϵ	ϵ																														
1	ϵ	ϵ	ϵ																														
	0	0.1	1																														
0	ϵ	1-8 ϵ	ϵ																														
0.1	ϵ	ϵ	ϵ																														
1	ϵ	ϵ	ϵ																														

where we suppose $0 < \epsilon < 0.01$. Here columns (rows) are labeled by outputs of agent 1 (2), and the entries of the matrix give probabilities for any pair of outputs. (It is easy to check that the condition on correlated probabilities in Section 11 is met.)

Thus the probabilities of production are correlated: when both agents work, an externality (“spillover” effect) comes into play, enabling each to produce a lot (i.e., 1 unit), albeit with low probability (i.e., 0.2); if only one

agent works, the benefits of this externality are removed, and he is only able to produce 0.1.

Disutility of effort is given by $d_1(0) = d_2(0) = 0$, $d_1(1) = 1$, $d_2(1) = 2$; and the utility of money by $u_1(x) = x^{1/2}$ and $u_2(x) = 2x^{1/2}$.

It is easy to check that, in order to implement (1, 1) as an SE, the bonus must be just high enough to satisfy $(\frac{1}{2} - \frac{9}{2}\epsilon)u_\alpha(B) \geq d_\alpha(1) - d_\alpha(0)$ for $\alpha \in \{1, 2\}$. For $B \geq 5$ the inequality becomes strict. Thus the SE-threshold bonus is less than 5 dollars; and, for $B = 5$, (1, 1) is the unique SE, as predicted by Theorem 3 (which still applies since $T = 1$). Thus the principal can induce both agents to work by putting up 5 dollars by way of a *competitive* prize.

Suppose, instead that he were to write an independent contract awarding x, y, z dollars¹³ for 0, 0.1, 1 units of output respectively (Since we deal only with non-discriminatory contracts, the reward does not depend on the name of the agent). If x, y, z are to induce both agents to work at effort level 1, we must have

$$(0.8)u_\alpha(x) + (0.2)u_\alpha(z) \geq d_\alpha(1)$$

and

$$(0.8)u_\alpha(x) + (0.2)u_\alpha(z) - d_\alpha(1) \geq u_\alpha(x) - d_\alpha(0) = u_\alpha(x)$$

These are the participation and incentive constraints respectively. It is clear that the latter implies that former. The incentive constraint may be rewritten $(0.2)[u_\alpha(z) - u_\alpha(x)] \geq d_\alpha(1)$, hence we have $u_\alpha(z) - u_\alpha(x) \geq 5d_\alpha(1)$. Since $u_\alpha(x) \geq 0$, it follows that $u_\alpha(z) \geq 5d_\alpha(1)$ which reduces to $z \geq 25$ for $\alpha \in \{1, 2\}$. Thus the principal's expected payout for the independent contract is $2[(0.8)x + (0.2)z] \geq (2)(0.2)z \geq (2)(0.2)(25) = 10$ dollars.

This is much higher than the payout of 5 dollars for the competitive prize.

It is worth noting that, in this example, the principal would have to pay out at least 10 dollars even if he were permitted to write *discriminatory* independent contracts $(x_\alpha, y_\alpha, z_\alpha)$ for $\alpha \in \{1, 2\}$. (Indeed the same argument shows that $z_\alpha \geq 25$ for $\alpha \in \{1, 2\}$.)

¹³For ease of computation we take x, y, z to be nonnegative. But in this example, competitive prizes are optimal even if we allow for "punishment" (via negative values of x, y, z).

13 Related Literature

We report on the model of Cowen & Glazer (1996), rephrasing it slightly. There is only one agent who undertakes countably infinite activities and chooses a uniform shirking rate S across all of them. The principal samples activities according to the Poisson distribution with mean M . Thus the probability that the principal samples H activities is $e^{-M} M^H / H!$. In effect M is his strategic variable and, by increasing M , he can scrutinize the agent more.

We may imagine that the agent tosses a coin (which comes heads with probability $1 - S$) independently for each activity, and works on that activity only if it has come up heads. Then the probability that he shirks on L out of the H activities sampled is $C(H, L) S^L (1 - S)^{H-L}$ (where, recall, $C(H, L) \equiv H! / (H - L)! L!$).

The agent receives the reward if the principal hears at least one message and if the fraction of messages reporting shirking is no more than some exogenously chosen critical value k . Thus the probability that the agent with shirking rate S receives the prize is

$$P(S) = \sum_{H=0}^{\infty} e^{-M} M^H / H! \sum_{L=0}^{int(Hk)} C(H, L) S^L (1 - S)^{H-L}.$$

where $int(Hk)$ is the largest integer that does not exceed Hk .

If $M > 0$ and $k > 0$, $P(S)$ is a downward sloping function of S . We reproduce the Figure 4 from Cowen & Glazer (1996).

The solid lines show $P(S)$ for $M = 3$ and $M = 4$. The dashed lines represent indifference curves of the agent, who needs to be compensated with a higher probability of the prize for shirking less. It *can* happen, for certain configurations of the agent's preferences, that he shirks more when the principal's scrutiny (i.e., M) goes up.

The agent's problem here may be thought of as his "best-reply" problem in an appropriately defined game in which his rival's strategy is held fixed in a particular way. Indeed, suppose that the rival works every day and produces exactly¹⁴ k units of output each day, for $0 < k < 1$. The agent himself produces exactly 0 if he shirks, and exactly 1 if he works. He chooses a "stationary mixed strategy" via a shirking rate. The principal samples days according to the Poisson distribution, and awards the prize to whoever produces more across the days in the sample. Finally the payoff of the agent depends on his shirking rate and the probability of winning the prize.

¹⁴Or, with probability close to 1.




Figure 4

The contrast with our model is clear. First, reduced scrutiny is *always* to the principal's advantage in our model, provided T is large; but in Cowen & Glazer (1996) this happens only *sometimes* for suitably chosen configurations of preferences as is evident from Figure 3. Next, the “opportunity curves” $P(S)$ for each agent in our model arise fundamentally from the *strategic competition* between them. Of course we could take a “partial equilibrium” approach, imagining the external rival agent as above. If we set $I_1 = \Omega$, the “best-reply” problem for the agent reduces to the model in Cowen & Glazer (1996). But the external agent's incentives remain unaccounted for in this partial equilibrium scenario.

14 Many Agents

Much of our analysis clearly continues to hold when there are more than two agents. In particular, for any fixed sample size n , there exist, as before,

finite thresholds

$$\begin{aligned}\tilde{B}(n) &\equiv \min\{B \in R_+ : \sigma^* \text{ is a DSE of } \Gamma(B, n)\} \\ B_*(n) &\equiv \inf\{B \in R_+ : \sigma^* \text{ is the unique SE of } \Gamma(B, n) \text{ for all } \tilde{B} > B\} \\ B(n) &\equiv \min\{B \in R_+ : \sigma^* \text{ is an SE of } \Gamma(B, n)\}\end{aligned}$$

and the corresponding optimal sample sizes will again typically be much smaller than T . We must have $B(n) \leq B_*(n) \leq \tilde{B}(n)$. Theorem 3 showed that, with two agents who are ignorant of each other's outputs, $B_*(n) = B(n)$ always (though, typically, $B(n) < \tilde{B}(n)$). We have not yet fully explored the relation between these three thresholds for multiple agents. But it is curious that, if $T = 1$ and if all agents can produce just two output levels and their probabilities of being productive are independent, then $\tilde{B}(n) = B_*(n) = B(n)$ (we get then a game of strategic substitutes as in Bulow et al (1985)).

15 General Contracts

A more (though not most) general formulation¹⁵ of a non-discriminatory contract, is in terms of a function

$$Z_+^T \times Z_+^T \xrightarrow{\psi} ([0, 1] \times R_+) \times ([0, 1] \times R_+)$$

Let $\psi(x, y) = (p, B, q, D)$. We understand the contract ψ to mean that if the principal observes outputs x, y by agents 1, 2 then he awards B, D dollars to agents 1, 2 with probability p, q respectively. Of course, we must require

$$\psi(x, y) = (p, B, q, D) \Rightarrow \psi(y, x) = (q, D, p, B)$$

to maintain the non-discriminatory character of ψ .

Such general contracts are not easily susceptible to calculations. But among them are contracts which operate as follows: the principal observes z_1, z_2 and awards $(p(z_1, z_2), B, 1 - p(z_1, z_2), B)$. They include our contracts (B, n) . Choosing n small is tantamount to making $p(z_1, z_2)$ large even when $\sum_{t=1}^T z_1(t) < \sum_{t=1}^T z_2(t)$, i.e., not always giving the bonus to the bigger producer, but deliberately "making a mistake" and giving it with probability $p(z_1, z_2)$ to the smaller producer to stimulate him to work hard.

¹⁵This line of inquiry, which we have not pursued here, was indicated by Jean-François Mertens.

16 Fine Information with Long Horizon: Further Considerations

Our analysis suggests that the principal could think of creating several prizes $(T_1, B_1, n_1), \dots, (T_k, B_k, n_k)$, breaking the T periods into k intervals of lengths T_1, \dots, T_k ; and award the prize (B_l, n_l) over the T_l days of the l th interval, wiping the past clean. This might be more efficient than a single prize (B, n) over the entire horizon, in that $\sum_{l=1}^k B_l < B$. Optimality would now determine an endogenous number of prizes and their timing.

We may also raise a different question here. Suppose agents move sequentially. In any period t , agent 1 first chooses his effort level $e_1 \in E_1$, after which chance moves and selects $q_1 \in Q_1$ with probability $p_1(q_1)$. Then agent 2, *knowing both e_1 and q_1* , picks e_2 from E_2 etc. We obtain an extensive form with perfect information. At its terminal nodes in $\Omega(T+1)$, it is useful to add a move of chance which embodies the principal's policy (B, n) . Chance selects every $\mathfrak{S} \subset T$, for $|\mathfrak{S}| = n$, with probability $1/C(T, n)$. Payoffs are determined at $\Omega(T+2)$ in the obvious manner.

In this extensive game, there exists a unique subgame perfect pure strategy equilibrium (unique, for generic specification of producible quantities q_α in Q_α and probability distribution $p_{\alpha,\omega}$ and disutilities $d_{\alpha,\omega}$ at each node ω in the tree) in which agents will *not* put in maximal effort everywhere (unless B is huge). Suppose the principal's budget B is of medium size and is *fixed*. For any sample size n in the above game $\Gamma(B, n)$, there is an expected total output $q(n)$ produced by the agents at its unique subgame perfect SE. We say that $\tilde{n}(T)$ is an optimal sample size if $q(\tilde{n}(T)) = \max\{q(n) : 1 \leq n \leq T\}$. We conjecture that $\tilde{n}(T)/T \rightarrow 0$ as $T \rightarrow \infty$ (for fixed B). This has been proved, under certain additional assumptions, in Dubey & Haimanko(2000).

17 Proofs

17.1 Preliminaries

If $\tilde{\omega} \equiv (\omega, e_1, e_2, q_1, q_2) \in \Omega(t+1)$ for $\omega \in \Omega(t)$, we say that $\tilde{\omega}$ is an immediate follower of ω . More generally, $\tilde{\omega}$ is a follower of ω if there exists a sequence $\omega_1, \omega_2, \dots, \omega_k$ of nodes in $\bar{\Omega} \equiv \coprod_{t=1}^{T+1} \Omega(t) = \Omega \coprod \Omega(T+1)$ such that $\omega = \omega_1$, $\tilde{\omega} = \omega_k$ and each ω_{l+1} is an immediate follower of ω_l for $1 \leq l \leq k-1$. Denote

$$F(\omega) = \{\tilde{\omega} \in \bar{\Omega} : \tilde{\omega} \text{ is a follower of } \omega\}.$$

Next let $\omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \in \Omega(t) \subset \Omega$. Define $Q(\alpha, \omega) \equiv$

$$\{((x_\alpha(\tau))_{\tau \in T \setminus \{t\}}, (y_\beta(\tau))_{\tau \in T}) \in Q_\alpha^{T \setminus \{t\}} \times Q_\beta^T : (x_\alpha(\tau), y_\beta(\tau))_{\tau=1}^{t-1} = (q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1}\}$$

i.e., $Q(\alpha, \omega)$ lists all possible outputs of both agents, across time, that are consistent with ω and that leave α 's output at ω unspecified. For any $(x, y) \in Q(\alpha, \omega)$ and $\mathfrak{S} \subset T$, with $t \in \mathfrak{S}$, put

$$\delta(x, y, \mathfrak{S}) \equiv \sum_{\tau \in \mathfrak{S}} y(\tau) - \sum_{\tau \in \mathfrak{S} \setminus \{t\}} x(\tau).$$

It is clear that α wins the bonus at ω , given \mathfrak{S} and (x, y) , if he produces more than $\delta(x, y, \mathfrak{S})$ at ω ; and ties for the bonus if he produces exactly $\delta(x, y, \mathfrak{S})$. Let

$$\overline{W}_\alpha = \{q \in Q_\alpha : q > \delta(x, y, \mathfrak{S})\}$$

$$W_\alpha = \{q \in Q_\alpha : q \geq \delta(x, y, \mathfrak{S})\}$$

$$\underline{W}_\alpha = \{q \in Q_\alpha : q = \delta(x, y, \mathfrak{S})\}$$

Then, conditional on the realization of $(x, y) \in Q(\alpha, \omega)$ and on the days in \mathfrak{S} being sampled, the effort $e \in E_\alpha$ at ω induces the probability $Prob_\alpha(x, y, e, \omega, \mathfrak{S})$ for α to win the bonus, where ¹⁶

$$\begin{aligned} Prob_\alpha(x, y, e, \omega, \mathfrak{S}) &= p_\alpha^e(\overline{W}_\alpha) + \frac{1}{2}p_\alpha^e(\underline{W}_\alpha) \\ &= \frac{1}{2}(p_\alpha^e(\overline{W}_\alpha) + p_\alpha^e(W_\alpha)) \end{aligned}$$

(The second equality comes from: $\underline{W}_\alpha = W_\alpha \setminus \overline{W}_\alpha$ and $\overline{W}_\alpha \subset W_\alpha$.) By assumption (II), $p_\alpha^{e^*}(V) \geq p_\alpha^e(V)$ for all $e \in E_\alpha$ and sets V that are of the type $\{q \in Q_\alpha : q \geq x\}$ or $\{q \in Q_\alpha : q > x\}$, x arbitrary; hence we immediately have

$$Prob_\alpha(x, y, e_\alpha^*, \omega, \mathfrak{S}) \geq Prob_\alpha(x, y, e, \omega, \mathfrak{S}) \quad (1)$$

for all $e \in E_\alpha$. Moreover, if $\underline{q}_\alpha \leq \delta(x, y, \mathfrak{S}) \leq \overline{q}_\alpha$ we have (again, by assumption (II) and the fact that $p_\alpha^{e^*}$ has full support)

$$Prob_\alpha(x, y, e_\alpha^*, \omega, \mathfrak{S}) > Prob_\alpha(x, y, e, \omega, \mathfrak{S}) \quad (2)$$

for all $e \in E_\alpha \setminus \{e_\alpha^*\}$

¹⁶For any $S \subset Q_\alpha$, $p_\alpha^e(S) \equiv \sum_{q \in S} p_\alpha^e(q)$

Next, suppose $\omega \equiv (e_1(\tau), e_2(\tau), q_1(\tau), q_2(\tau))_{\tau=1}^T \in \Omega(T+1)$ is a terminal node and $\mathfrak{S} \subset T$. Define

$$Prob_\alpha(\omega, \mathfrak{S}) = \begin{cases} 1 & \text{if } \sum_{\tau \in \mathfrak{S}} q_\alpha(\tau) > \sum_{\tau \in \mathfrak{S}} q_\beta(\tau) \\ 1/2 & \text{if } \sum_{\tau \in \mathfrak{S}} q_\alpha(\tau) = \sum_{\tau \in \mathfrak{S}} q_\beta(\tau) \\ 0 & \text{if } \sum_{\tau \in \mathfrak{S}} q_\alpha(\tau) < \sum_{\tau \in \mathfrak{S}} q_\beta(\tau) \end{cases}$$

Consider two arbitrary maps¹⁷ $\sigma_\alpha, \sigma_\beta$ from Ω to E_α, E_β respectively. The probability, induced by $(\sigma_\alpha, \sigma_\beta)$, that the play of the game goes through ω and α wins the bonus under sample \mathfrak{S} is given by:

$$Prob_\alpha((\sigma_\alpha, \sigma_\beta)|\mathfrak{S}, \omega) = \sum_{\tilde{\omega} \in F(\omega) \cap \Omega(T+1)} p^{(\sigma_\alpha, \sigma_\beta)}(\tilde{\omega}) Prob_\alpha(\tilde{\omega}, \mathfrak{S})$$

When \mathfrak{S} ranges over all of \mathcal{C}_n and all samples in \mathcal{C}_n are picked with uniform probability, this event has probability

$$Prob_\alpha((\sigma_\alpha, \sigma_\beta)|n, \omega) = \sum_{\mathfrak{S} \in \mathcal{C}_n} P(n) Prob_\alpha((\sigma_\alpha, \sigma_\beta)|\mathfrak{S}, \omega) \quad (3)$$

Finally, the overall probability that α wins the bonus in the game, under $(\sigma_\alpha, \sigma_\beta)$ and sample size n , is

$$Prob_\alpha^n(\sigma_\alpha, \sigma_\beta) = Prob_\alpha((\sigma_\alpha, \sigma_\beta)|n, \omega^*)$$

It is obvious that

$$Prob_\alpha^n(\sigma_\alpha, \sigma_\beta) = \sum_{\omega \in \Omega(t)} Prob_\alpha((\sigma_\alpha, \sigma_\beta)|n, \omega) \quad (4)$$

for every $1 \leq t \leq T+1$.

We now state two useful lemmas. First a

Definition Fix $\omega \in \Omega$. Suppose $\hat{\sigma}_\alpha$ and $\tilde{\sigma}_\alpha$ are two maps from Ω to E_α which satisfy: (i) $\hat{\sigma}_\alpha(\tilde{\omega}) = \tilde{\sigma}_\alpha(\tilde{\omega})$ for $\tilde{\omega} \in \Omega \setminus \{\omega\}$; (ii) $\hat{\sigma}_\alpha(\tilde{\omega}) = \tilde{\sigma}_\alpha(\tilde{\omega}) = e_\alpha^*$ for $\tilde{\omega} \in F(\omega)$, (iii) $\hat{\sigma}_\alpha(\omega) = e_\alpha^* \neq \tilde{\sigma}_\alpha(\omega)$. Then we write: $\hat{\sigma}_\alpha \succ_\omega \tilde{\sigma}_\alpha$.

Lemma 1 Fix $\omega \in \Omega(t)$ and $\mathfrak{S} \subset T$. Suppose $\hat{\sigma}_\alpha \succ_\omega \tilde{\sigma}_\alpha$. Then

- (a) $Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*)|\mathfrak{S}, \omega) \geq Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*)|\mathfrak{S}, \omega)$
- (b) $Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*)|n, \omega) \geq Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*)|n, \omega)$
- (c) $Prob_\alpha^n(\hat{\sigma}_\alpha, \sigma_\beta^*) \geq Prob_\alpha^n(\tilde{\sigma}_\alpha, \sigma_\beta^*)$

¹⁷not necessarily strategies in $\Sigma_\alpha, \Sigma_\beta$.

Proof

It is evident that, if $t \in \mathfrak{S}$,

$$\begin{aligned}
 & Prob_\alpha((\sigma_\alpha, \sigma_\beta^*) | \mathfrak{S}, \omega) = p^{(\sigma_\alpha, \sigma_\beta^*)}(\omega) \times \\
 & \left(\sum_{(x,y) \in Q(\alpha, \omega)} \prod_{\tau=t}^T p_\beta^{e_\beta^*}(y(\tau)) \right) \left(\prod_{\tau=t+1}^T p_\alpha^{e_\alpha^*}(x(\tau)) \right) Prob_\alpha(x, y, \sigma_\alpha(\omega), \omega, \mathfrak{S}) \quad (5)
 \end{aligned}$$

for all σ_α , in particular $\sigma_\alpha = \hat{\sigma}_\alpha$ or $\tilde{\sigma}_\alpha$. By (1) and (5), and the fact that $p^{(\hat{\sigma}_\alpha, \sigma_\beta^*)}(\omega) = p^{(\tilde{\sigma}_\alpha, \sigma_\beta^*)}(\omega)$ on account of (i) in the definition of “ \succ_ω ”, we get $Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}, \omega) \geq Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}, \omega)$. If t is not in \mathfrak{S} , it is obvious that $Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}, \omega) = Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}, \omega)$. This proves (a) of Lemma 1. Now (b) is obvious from (a) and (3). Finally (c) is obvious from (4), and (b), and the obvious fact that $Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*) | n, \omega') = Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*) | n, \omega')$ for all $\omega' \in \Omega(t) \setminus \{\omega\}$.

Lemma 2 For any map $\hat{\sigma}_\alpha : \Omega \rightarrow E_\alpha$, and $1 \leq n \leq T$:

$$Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) \geq Prob_\alpha^n(\hat{\sigma}_\alpha, \sigma_\beta^*)$$

Proof

Introduce a total order \succ on all the nodes in Ω such that, if $\omega \in \Omega(t)$ and $\tilde{\omega} \in \Omega(\tilde{t})$ and $\tilde{t} > t$, then $\tilde{\omega} \succ \omega$. Write $\Omega \equiv \{\omega_1, \omega_2, \dots, \omega_L\}$ with $\omega_1 \succ \omega_2 \succ \dots \succ \omega_L$. Go from σ_α to σ_α^* through a sequence of L transitions:

$$\sigma_\alpha \equiv \sigma_\alpha^1 \longrightarrow \sigma_\alpha^2 \longrightarrow \dots \longrightarrow \sigma_\alpha^{L+1} \equiv \sigma_\alpha^*$$

In the l^{th} transition, change $\sigma_\alpha^{l-1}(\omega_l)$ to e_α^* and leave the rest of the choices according to σ_α^{l-1} ; denote the resulting map by σ_α^l . Clearly $\sigma_\alpha^{l+1} \succ_{\omega_l} \sigma_\alpha^l$.

By (c) of Lemma 1,

$$Prob_\alpha^n(\sigma_\alpha^{l+1}, \sigma_\beta^*) \geq Prob_\alpha^n(\sigma_\alpha^l, \sigma_\beta^*)$$

for $1 \leq l \leq L$. Therefore

$$Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) \geq Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*)$$

proving Lemma 2.

17.2 Proof of Theorem 4

17.2.1 Motivation and Definition of f_1, f_2

As we said, $f_\alpha(\omega)$ is the largest sample size which still leaves some incentive for agent α to put in maximal effort e_α^* at the node ω , *when he knows $\omega \in \Omega$ perfectly.*

Let us consider what could destroy this incentive. Suppose α knows at ω that his rival has already produced much more than him in the days sampled in the past.....in fact so much more, that even if he produces his maximum \bar{q}_α and his rival β produces his minimum \underline{q}_β on all future days remaining in the sample, he still cannot equal or beat β . In this scenario, it would make no sense for α to work at ω . We call this *the bad scenario for α at ω , given sample size n .* The good scenario is when his rival is in a bad scenario, i.e., α is so far ahead of β that he will win without working even if β turns lucky and produces \bar{q}_β on all future days while he himself produces only \underline{q}_α . In the good scenario, too, there is no incentive for α to work.

Notice that these scenarios can only arise if two things happen at once: α knows a lot of the past history (i.e., has refined information) *and* sufficiently many days of the past are sampled (i.e., the sample size n is not too small).

With this motivation in mind, let us now do a precise calculation. Suppose the history of outputs that leads to $\omega \in \Omega(t)$ is $(q_\alpha(t), q_\beta(t))_{\tau=1}^{t-1}$ and suppose that the sample size is n . When will a bad scenario be inescapable for α at ω ? If $n \leq T - t + 1$ there is a *positive* probability that all the n days in the sample will consist of ω and $n - 1$ days after ω , and that α 's rival β will produce exactly \underline{q}_β on these n days. By putting in effort e_α^* at ω , α can boost¹⁸ the probability of producing \bar{q}_α at ω , without lowering his chances of winning the bonus in any other event. So the bad scenario is avoided. Now suppose $n > T - t + 1$. Let us consider an optimal way of avoiding the bad scenario for α . To this end, first let ω and all the $T - t$ days of the future be sampled. This still leaves $r \equiv n - (T - t + 1) = n - T + t - 1$ more days to be sampled from the past of ω . Denote $\Delta(\tau) \equiv q_\alpha(\tau) - q_\beta(\tau)$ for $1 \leq \tau < t - 1$. For any integer $-\infty < j \leq t - 1$, define

$$\bar{h}(j, \omega) = \begin{cases} \max\{\sum_{\tau \in \tilde{\mathfrak{S}}} \Delta(\tau) : \tilde{\mathfrak{S}} \subset \{1, \dots, t - 1\}, |\tilde{\mathfrak{S}}| = j\} & \text{if } j > 0 \\ 0 & \text{if } j \leq 0 \end{cases}$$

The subset \mathfrak{S} of days that could be sampled from the past of ω , to best help α avoid the bad scenario, would be a \mathfrak{S} that achieves $\bar{h}(r, \omega)$ and maximizes α 's lead over β . Denote integers by \mathbf{Z} and define

¹⁸By Assumption (I)

$$\underline{f}_\alpha(\omega) = \max\{k \in \mathbf{Z} : 1 \leq k \leq T, (T-t+1)(\bar{q}_\alpha - \underline{q}_\beta) + \bar{h}(k-T+t-1, \omega) \geq 0\}$$

It is easy to check that a bad scenario can be avoided if, and only if, $n \leq \underline{f}_\alpha(\omega)$. Arguing in the same fashion, a good scenario for α can be avoided if, and only if, $n \leq \bar{f}_\alpha(\omega)$, where

$$\bar{f}_\alpha(\omega) = \max\{k \in \mathbf{Z} : 1 \leq k \leq T, (T-t+1)(\underline{q}_\alpha - \bar{q}_\beta) + \underline{h}(k-T+t-1, \omega) \leq 0\}$$

and $\underline{h}(j, \omega)$ is defined exactly as $\bar{h}(j, \omega)$ replacing “max” by “min”. Since we must avoid both these scenarios, we put

$$f_\alpha(\omega) = \min\{\underline{f}_\alpha(\omega), \bar{f}_\alpha(\omega)\}.$$

This completes the definition of the functions¹⁹ f_α for $\alpha \in \{1, 2\}$, and thereby of g . Recall

$$g(\Psi) = \min_{\alpha \in \{1, 2\}} \min_{S \in I_\alpha} \max\{f_\alpha(\omega) : \omega \in S\}$$

17.2.2 Lemmas

Lemma 3 Let $a_1, \dots, a_K, \bar{a}, \underline{a}$ be real numbers such that $a_1 \leq \bar{a}$, $a_K \geq \underline{a}$ and $|a_{l+1} - a_l| \leq \bar{a} - \underline{a}$ for $1 \leq l \leq K - 1$. Then $\underline{a} \leq a_j \leq \bar{a}$ for some $1 \leq j \leq K$.

Proof Obvious.

Lemma 4 Assume $n \leq f_\alpha(\omega)$ for $\omega \in \Omega(t)$ and $\alpha \in \{1, 2\}$. Then there exists $\mathfrak{S} \in \mathcal{C}_n$ and $(x, y) \in Q(\alpha, \omega)$ such that: $t \in \mathfrak{S}$ and $\underline{q}_\alpha \leq \delta(x, y, \mathfrak{S}) \leq \bar{q}_\alpha$.

Proof Let $\omega \equiv (e_\alpha(\tau), e_\beta(\tau), q_\alpha(\tau), q_\beta(\tau))_{\tau=1}^{t-1} \in \Omega(t)$. Denote $T^* \equiv T \setminus \{t\}$. We shall say that $(x_\alpha(\tau), y_\beta(\tau))_{\tau \in T^*}$ is consistent with ω if $x_\alpha(\tau) = q_\alpha(\tau)$ and $y_\beta(\tau) = q_\beta(\tau)$ for $1 \leq \tau \leq t - 1$; and $x_\alpha(\tau) \in Q_\alpha$ and $y_\beta(\tau) \in Q_\beta$ for $\tau \in T^*$. Since $f_\alpha(\omega) = \min\{\bar{f}_\alpha(\omega), \underline{f}_\alpha(\omega)\}$, $n \leq \bar{f}_\alpha(\omega)$ and $n \leq \underline{f}_\alpha(\omega)$.

¹⁹With *two* agents $f_1(\omega) = f_2(\omega)$. But the functions $\underline{f}_\alpha(\omega)$, $\bar{f}_\alpha(\omega)$ can be defined, with obvious modifications, for the case of N agents. Then, if we define g as above, with “ $\min_{\alpha \in \{1, 2\}}$ ” replaced by “ $\min_{\alpha \in \{1, \dots, N\}}$ ”, Theorem 4 remains intact (with the same proof).

Now, $n \leq \bar{f}_\alpha(\omega)$ implies²⁰

$\exists \bar{\mathfrak{S}} \in \mathcal{C}_n$ and $(\bar{x}_\alpha(\tau), \bar{y}_\beta(\tau))_{\tau \in T^*}$ consistent with ω such that:

$$t \in \bar{\mathfrak{S}} \text{ and } \sum_{\tau \in \bar{\mathfrak{S}} \setminus \{t\}} (\bar{x}_\alpha(\tau) - \bar{y}_\beta(\tau)) \leq \bar{q}_\beta - \underline{q}_\alpha \quad (6)$$

i.e., ω is not a good scenario for α (he cannot clearly win if he produces \underline{q}_α and his rival produces \bar{q}_β at ω); similarly, $n \leq \underline{f}_\alpha(\omega)$ implies

$\exists \underline{\mathfrak{S}} \in \mathcal{C}_n$ and $(\underline{x}_\alpha(\tau), \underline{y}_\beta(\tau))_{\tau \in T^*}$ consistent with ω such that:

$$t \in \underline{\mathfrak{S}} \text{ and } \sum_{\tau \in \underline{\mathfrak{S}} \setminus \{t\}} (\underline{x}_\alpha(\tau) - \underline{y}_\beta(\tau)) \geq \underline{q}_\beta - \bar{q}_\alpha \quad (7)$$

i.e., ω is not a bad scenario for α (he cannot clearly lose if he produces \bar{q}_α and his rival produces \underline{q}_β at ω). Consider any sequence of samples $\bar{\mathfrak{S}} \equiv \mathfrak{S}_1 \rightarrow \mathfrak{S}_2 \rightarrow \dots \rightarrow \mathfrak{S}_k \equiv \underline{\mathfrak{S}}$ such that, in each transition $\mathfrak{S}_l \rightarrow \mathfrak{S}_{l+1}$, one day $\bar{\tau}_l \in \bar{\mathfrak{S}}$ is removed from \mathfrak{S}_l and is replaced by a day $\underline{\tau}_l \in \underline{\mathfrak{S}}$ (which is not necessarily distinct from $\bar{\tau}_l$), making sure that $\mathfrak{S}_{l+1} \in \mathcal{C}_n$. Notice $t \in \mathfrak{S}_l$ for $1 \leq l \leq k$. For each such l , define $(x_\alpha^l(\tau), y_\beta^l(\tau))_{\tau \in T^*}$, consistent with ω , by

$$(x_\alpha^l(\tau), y_\beta^l(\tau)) = \begin{cases} (\bar{x}_\alpha(\tau), \bar{y}_\beta(\tau)) & \text{if } \tau \in T^* \setminus \{\bar{\tau}_1, \dots, \bar{\tau}_{l-1}\} \\ (\underline{x}_\alpha(\tau), \underline{y}_\beta(\tau)) & \text{otherwise} \end{cases}$$

Consider $\Delta(l) \equiv \sum_{\tau \in \mathfrak{S}_l \setminus \{t\}} x_\alpha^l(\tau) - y_\beta^l(\tau)$. Notice $\Delta(1) = \sum_{\tau \in \bar{\mathfrak{S}} \setminus \{t\}} \bar{x}_\alpha(\tau) - \bar{y}_\beta(\tau) \leq \bar{q}_\beta - \underline{q}_\alpha$ and $\Delta(k) = \sum_{\tau \in \underline{\mathfrak{S}} \setminus \{t\}} \underline{x}_\alpha(\tau) - \underline{y}_\beta(\tau) \geq \underline{q}_\beta - \bar{q}_\alpha$. Moreover, $|\Delta(l+1) - \Delta(l)| = |(\underline{x}_\alpha(\underline{\tau}_l) - \underline{y}_\beta(\underline{\tau}_l)) - (\bar{x}_\alpha(\bar{\tau}_l) - \bar{y}_\beta(\bar{\tau}_l))| \leq (\bar{q}_\alpha - \underline{q}_\beta) - (\underline{q}_\alpha - \bar{q}_\beta) = (\bar{q}_\beta - \underline{q}_\alpha) - (\underline{q}_\beta - \bar{q}_\alpha)$. Apply Lemma 3 (with $\{a_1, \dots, a_K\} \equiv \{\Delta(1), \dots, \Delta(k)\}$, $\bar{a} \equiv \bar{q}_\beta - \underline{q}_\alpha$, $\underline{a} \equiv \underline{q}_\beta - \bar{q}_\alpha$). We conclude that there exists $\mathfrak{S} \in \mathcal{C}_n$ and $(x_\alpha(\tau), y_\beta(\tau))_{\tau \in T^*}$, consistent with ω , such that $t \in \mathfrak{S}$ and $\underline{q}_\beta - \bar{q}_\alpha \leq \sum_{\tau \in \mathfrak{S} \setminus \{t\}} (x_\alpha(\tau) - y_\beta(\tau)) \leq \bar{q}_\beta - \underline{q}_\alpha$. Rewrite this inequality as

$$\bar{q}_\beta \geq \underline{q}_\alpha + \sum_{\tau \in \mathfrak{S} \setminus \{t\}} (x_\alpha(\tau) - y_\beta(\tau)) \equiv \underline{A}$$

$$\underline{q}_\beta \leq \bar{q}_\alpha + \sum_{\tau \in \mathfrak{S} \setminus \{t\}} (x_\alpha(\tau) - y_\beta(\tau)) \equiv \bar{A}$$

²⁰(For $\tau \in \bar{\mathfrak{S}}$, $\tau > t$ we may in particular take $\bar{x}_\alpha(\tau) = \underline{q}_\alpha$, $\bar{y}_\beta(\tau) = \bar{q}_\beta$; and for $\tau \in \underline{\mathfrak{S}}$, $\tau > t$ take $\underline{x}_\alpha(\tau) = \bar{q}_\alpha$, $\underline{y}_\beta(\tau) = \underline{q}_\beta$)

Again apply Lemma 3 (with $\underline{a} \equiv \underline{A}$, $\bar{a} \equiv \bar{A}$, $\{a_1, \dots, a_K\} \equiv \{\underline{q}_\beta \equiv q_\beta^1, q_\beta^2, \dots, q_\beta^{m(\beta)} \equiv \bar{q}_\beta\}$) to conclude that there exists $q_\beta^l \in Q_\beta$ such that $\underline{A} \leq q_\beta^l \leq \bar{A}$. Put $y_\beta(t) = q_\beta^l$. This defines $y \in Q_\beta^T$ (since only the t^{th} -component of y was unspecified) and clearly $(x, y) \in Q(\alpha, \omega)$. Then x, y, \mathfrak{S} accomplish the requirement of Lemma 4.

17.2.3 Completion of the Proof of Theorem 4

We claim that

$$B(\Psi, n) < \infty \Leftrightarrow \text{Prob}_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > \text{Prob}_\alpha^n(\sigma_\alpha, \sigma_\beta^*) \quad (8)$$

for $\alpha \in \{1, 2\}$ and $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$. First we show \Rightarrow . Suppose $\text{Prob}_\alpha^n(\sigma_\alpha, \sigma_\beta^*) \geq \text{Prob}_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*)$ for some $\sigma_\alpha \neq \sigma_\alpha^*$. For any $\omega \in \Omega(T+1)$, denote $\omega \equiv (e_\alpha^\omega(\tau), e_\beta^\omega(\tau), q_\alpha^\omega(\tau), q_\beta^\omega(\tau))_{\tau=1}^T$. Now $p^{(\sigma_\alpha^*, \sigma_\beta^*)}(\omega) > 0$ for $\omega \in \Omega(T+1)$ implies $e_\alpha^\omega(\tau) = e_\alpha^*$ for $\tau = 1, \dots, T$. Hence $\sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha^*, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T) = d_\alpha(e_\alpha^*, \dots, e_\alpha^*)$. On the other hand (by $\sigma_\alpha \neq \sigma_\alpha^*$, by Assumption (I) and by the fact that the p_α^e and p_β^e have full support) there exists $\omega \in \Omega(T+1)$ such that $p^{(\sigma_\alpha, \sigma_\beta^*)}(\omega) > 0$ and $e_\alpha^\omega(\tau) \neq e_\alpha^*$ for some τ . By Assumption (IV), we deduce $\sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T) > \sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha^*, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T)$. Then α could switch from σ_α^* to σ_α , saving disutility and still not losing on the reward, contradicting that $(\sigma_\alpha^*, \sigma_\beta^*)$ is an SE of $\Gamma(\Psi, B, n)$, for any arbitrary B . But then $B(\Psi, n) = \infty$, a contradiction, proving \Rightarrow . Next we show \Leftarrow . Suppose $\text{Prob}_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > \text{Prob}_\alpha^n(\sigma_\alpha, \sigma_\beta^*)$ for all $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$. Let

$$\delta = \min_{\alpha \in \{1, 2\}} \min_{\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}} (\text{Prob}_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) - \text{Prob}_\alpha^n(\sigma_\alpha, \sigma_\beta^*)).$$

Since Σ_α is finite, $\delta > 0$. By Assumption (III), there exists a finite B such that $\delta u_\alpha(B) \geq d_\alpha((e_\alpha^*, \dots, e_\alpha^*)) - d_\alpha(e)$ for $\alpha \in \{1, 2\}$ and $e \in E_\alpha^T \setminus \{(e_\alpha^*, \dots, e_\alpha^*)\}$. But then $(\sigma_\alpha^*, \sigma_\beta^*)$ must be an SE of $\Gamma(\Psi, B, n)$, showing that $B(\Psi, n) \leq B < \infty$. This proves \Leftarrow and thereby (8).

Thus, to prove theorem 4, it suffices to show

$$n \leq g(\Psi) \Leftrightarrow \text{Prob}_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > \text{Prob}_\alpha^n(\sigma_\alpha, \sigma_\beta^*) \quad (9)$$

for $\alpha \in \{1, 2\}$ and $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$.

Suppose $n > g(\Psi)$. Then there exists an agent α and an information set $\tilde{S} \in I_\alpha$ where $n > f(\tilde{\omega})$ for all $\tilde{\omega} \in \tilde{S}$. By the definition of f , any node in \tilde{S} is either a good scenario for α (i.e., $n > \bar{f}_\alpha(\tilde{\omega})$) or a bad scenario for α (i.e.,

$n > \underline{f}_\alpha(\omega)$). Let \tilde{S} correspond to $(\tilde{e}_\alpha(\tau))_{\tau=1}^{t-1}$ and $K \in J_\alpha(t)$ in accordance with (ii) of Assumption (V). Consider S that corresponds to $(\tilde{e}_\alpha^*(\tau))_{\tau=1}^{t-1}$ and K . Then (using Assumption (I)) some node in S is reached with positive probability under $(\sigma_\alpha^*, \sigma_\beta^*)$. Since the definition of $f_\alpha(\omega)$ depends only on the history of *outputs* that lead to ω , each node in S is also either a good scenario or a bad scenario. (Indeed output histories over S and \tilde{S} both yield K .) Therefore at *all* terminal nodes that follow from an arbitrary $\omega \in S$, α wins (loses) the bonus if ω is a good (bad) scenario, no matter what moves the agents make after ω . Let α change from σ_α^* to σ_α where $\sigma_\alpha(\omega) \neq e_\alpha^*$ for all $\omega \in S$, and σ_α is the same as σ_α^* at all other information sets. Then $Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*) = Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*)$. This proves “ \Leftarrow ” of (9).

Now suppose $n \leq g(\Psi)$. Consider any $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$. Let

$$t = \max\{\tau : 1 \leq \tau \leq T, \sigma_\alpha(\omega) \neq e_\alpha^* \text{ for some } \tilde{\omega} \in \Omega(\tau)\}.$$

Then there exists $S \in I_\alpha$ such that $\emptyset \neq S \subset \Omega(t)$, $\sigma_\alpha(\tilde{\omega}) \neq e_\alpha^*$ for all $\tilde{\omega} \in S$ and $\sigma_\alpha(\tilde{\omega}) = e_\alpha^*$ for all $\tilde{\omega} \in \cup_{\tau=t+1}^T \Omega(\tau)$. First notice that there is some node $\omega \in S$ which is relevant for σ_α (otherwise σ_α is equivalent to σ_α^*). By perfect recall (see (i) of Assumption (V)) all nodes in S are relevant for σ_α . By the definition of g , there is a node $\hat{\omega} \in S$ with $n \leq f_\alpha(\hat{\omega})$. By Assumption (I) and that fact that p_α and p_β have full support, there is a node $\omega' \in S$ which is reached with positive probability and which has the same history of outputs as $\hat{\omega}$. Hence $n \leq f_\alpha(\omega')$ also, since f_α depends only on the history of outputs.

By Lemma 4, there exist $(x', y') \in Q(\alpha, \omega')$ and $\mathfrak{S}' \in \mathcal{C}_n$ such that: $t' \in \mathfrak{S}'$ and $\underline{q}_\alpha \leq \delta(x', y', \mathfrak{S}') \leq \overline{q}_\alpha$. Consider any two maps $\hat{\sigma}_\alpha$ and $\tilde{\sigma}_\alpha$ from Ω to E_α (not necessarily strategies in Σ_α) such that $\hat{\sigma}_\alpha \succ_{\omega'} \tilde{\sigma}_\alpha$. By (2), (3) and (5) we have

$$Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}', \omega') > Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*) | \mathfrak{S}', \omega') \quad (10)$$

Then, by (a) of Lemma 1 and (10), and summing over $\mathfrak{S} \in \mathcal{C}_n$ as in (3), we obtain

$$Prob_\alpha((\hat{\sigma}_\alpha, \sigma_\beta^*) | n, \omega') > Prob_\alpha((\tilde{\sigma}_\alpha, \sigma_\beta^*) | n, \omega') \quad (11)$$

Go from σ_α to σ_α^* through the sequence of transitions $\sigma_\alpha \equiv \sigma_\alpha^1 \rightarrow \dots \rightarrow \sigma_\alpha^{L+1} \equiv \sigma_\alpha^*$ as in the proof of Lemma 2. By (c) of Lemma 1,

$$Prob_\alpha^n(\sigma_\alpha^{l+1}, \sigma_\beta^*) \geq Prob_\alpha^n(\sigma_\alpha^l, \sigma_\beta^*)$$

for all l . But, by (11) and (4), we have strict inequality at least once (at the transition involving ω'), proving

$$Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*)$$

Since σ_α was arbitrary element of $\Sigma_\alpha \setminus \{\sigma_\alpha^*\}$, this completes the proof.

17.3 Proof of Theorem 1

Consider any $\omega \in \Omega(T+1)$. Recall $\omega \equiv (e_\alpha^\omega(\tau), e_\beta^\omega(\tau), q_\alpha^\omega(\tau), q_\beta^\omega(\tau))_{\tau=1}^T$. Suppose $(\sigma_\alpha^*, \sigma_\beta^*)$ is an SE of $\Gamma(\Psi, B(\Psi, n), n)$. Then we have

$$\begin{aligned} Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) u_\alpha(B(\Psi, n)) - \sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha^*, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T) &\geq \\ Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*) u_\alpha(B(\Psi, n)) - \sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T) & \quad (12) \end{aligned}$$

for any $\sigma \in \Sigma \setminus \{\sigma_\alpha^*\}$.

Now, as argued in 17.2.3,

$$\sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha^*, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T) > \sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha, \sigma_\beta^*)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T)$$

Then by (12), $Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*)$. By assumption (II), $B > B(\Psi, n)$ implies $u_\alpha(B) > u_\alpha(B(\Psi, n))$ and so all the inequalities in (12) become strict when we replace $B(\Psi, n)$ by B . This shows that $(\sigma_\alpha^*, \sigma_\beta^*)$ is a strict SE of $\Gamma(\Psi, B, n)$, proving the “if” part of (i).

Next suppose that $(\sigma_\alpha^*, \sigma_\beta^*)$ is a strict SE of $\Gamma(\Psi, B, n)$. Then all the inequalities of (12) must be strict. By assumption (II), they will remain strict if we reduce B slightly to $B - \epsilon$, showing that $(\sigma_\alpha^*, \sigma_\beta^*)$ is an SE of $\Gamma(\Psi, B - \epsilon, n)$. We conclude that $B > B(\Psi, n)$.

The proof of (ii) is identical to the proof of (i), replacing σ_β^* by an arbitrary σ_β everywhere.

17.4 Proof of Theorem 2

First we prove two lemmas.

Lemma 5 Let $n \leq (\Delta/\tilde{\Delta}) + 1$ and $(x(1), \dots, x(n-1)) \in Q_\alpha^{n-1}$. There exists $(y(1), \dots, y(n)) \in Q_\beta^n$ such that $\sum_{\tau=1}^{n-1} x(\tau) + \underline{q}_\alpha \leq \sum_{\tau=1}^n y(\tau) \leq \sum_{\tau=1}^{n-1} x(\tau) + \bar{q}_\alpha$.

Proof

Consider the set $Y \equiv \{\sum_{\tau=1}^n y(\tau) : y(\tau) \in Q_\beta\}$, and write $Y = \{z^1, \dots, z^{nm(\beta)}\}$ with $z^1 < \dots < z^{nm(\beta)}$. Note $z^1 = n\underline{q}_\beta$ and $z^{nm(\beta)} = n\bar{q}_\beta$. Now, from $n \leq (\Delta/\tilde{\Delta}) + 1$, we have

$$\sum_{\tau=1}^{n-1} x(\tau) + \underline{q}_\alpha \leq (n-1)\bar{q}_\alpha + \underline{q}_\alpha \leq n\bar{q}_\beta$$

and

$$\sum_{\tau=1}^{n-1} x(\tau) + \bar{q}_\alpha \geq (n-1)\underline{q}_\alpha + \bar{q}_\alpha \geq n\underline{q}_\beta$$

Then, by Lemma 3 (with $\bar{a} = \sum_{\tau=1}^{n-1} x(\tau) + \bar{q}_\alpha$, $\underline{a} = \sum_{\tau=1}^{n-1} x(\tau) + \underline{q}_\alpha$ and $\{a_1, \dots, a_K\} = \{z^1, \dots, z^{nm(\beta)}\}$), there exists $(y(1), \dots, y(n)) \in Q_\beta^n$ such that $\sum_{\tau=1}^{n-1} x(\tau) + \underline{q}_\alpha \leq \sum_{\tau=1}^n y(\tau) \leq \sum_{\tau=1}^{n-1} x(\tau) + \bar{q}_\alpha$.

Lemma 6 Suppose neither agent can observe his rival in Ψ and $n \leq (\Delta/\tilde{\Delta}) + 1$. Fix $\omega \in \Omega(t)$, $\alpha \in \{1, 2\}$ and $\sigma_\beta \in \Sigma_\beta$. Let $\hat{\sigma}_\alpha$ and $\tilde{\sigma}_\alpha$ be two maps from Ω to E_α such that $\hat{\sigma}_\alpha \succ_\omega \tilde{\sigma}_\alpha$. Also assume that $p^{(\hat{\sigma}_\alpha, \sigma_\beta)}(\omega) > 0$. Then

$$Prob_\alpha^n(\hat{\sigma}_\alpha, \sigma_\beta) > Prob_\alpha^n(\tilde{\sigma}_\alpha, \sigma_\beta)$$

Proof Let $\tilde{\omega}$ be any node in $\Omega(t)$. Since β cannot observe α , the probability distribution on output vectors that β can produce after $\tilde{\omega}$ depends only on σ_β and $\tilde{\omega}$. Denote it by $Pr^{\sigma_\beta, \tilde{\omega}}$ and note that its support is $(Q_\beta)^{T-t+1}$. Therefore, if $t \in \tilde{\mathfrak{S}} \in \mathcal{C}_n$ and $\sigma_\alpha = \hat{\sigma}_\alpha$ or $\tilde{\sigma}_\alpha$, we have

$$Prob_\alpha((\sigma_\alpha, \sigma_\beta)|\tilde{\mathfrak{S}}, \tilde{\omega}) = p^{(\sigma_\alpha, \sigma_\beta)}(\tilde{\omega}) \times \left[\sum_{(x, y) \in Q(\alpha, \tilde{\omega})} (Pr^{\sigma_\beta, \tilde{\omega}}((y(\tau))_{\tau=t}^T)) \right. \\ \left. \left(\prod_{\tau=t+1}^T p_\alpha^{e_\alpha^*}(x(\tau)) \right) Prob_\alpha(x, y, \sigma_\alpha(\tilde{\omega}), \tilde{\omega}, \tilde{\mathfrak{S}}) \right] \quad (13)$$

Now $Prob_\alpha^n(\sigma_\alpha, \sigma_\beta) = \sum_{\tilde{\omega} \in \Omega(t)} \sum_{\tilde{\mathfrak{S}} \in \mathcal{C}_n} P(n) Prob_\alpha((\sigma_\alpha, \sigma_\beta)|\tilde{\mathfrak{S}}, \tilde{\omega})$. By Lemma 5 and (3), $Prob_\alpha(x, y, \hat{\sigma}_\alpha(\omega), \omega, \mathfrak{S}) > Prob_\alpha(x, y, \tilde{\sigma}_\alpha(\omega), \omega, \mathfrak{S})$; and by (2) we have “ \geq ” for $\tilde{\omega} \in \Omega(t) \setminus \{\omega\}$. Moreover it is clear that $p^{(\sigma_\alpha, \sigma_\beta)}(\tilde{\omega})$ is independent of $\sigma_\alpha = \hat{\sigma}_\alpha$ or $\tilde{\sigma}_\alpha$ for $\tilde{\omega} \in \Omega(t)$ given the condition (i) in the definition of \succ_ω . Thus, using (13) and (4),

$$Prob_\alpha^n(\hat{\sigma}_\alpha, \sigma_\beta) > Prob_\alpha^n(\tilde{\sigma}_\alpha, \sigma_\beta)$$

17.4.1 Completion of the Proof of Theorem 2

Consider any $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$. Let

$$t = \max\{\tau : 1 \leq \tau \leq T, \sigma_\alpha(\omega) \neq e_\alpha^* \text{ for some } \tilde{\omega} \in \Omega(\tau)\}.$$

Then there exists $S \in I_\alpha$ such that $S \cap \Omega(t) \neq \emptyset$, $\sigma_\alpha(\tilde{\omega}) \neq e_\alpha^*$ for all $\tilde{\omega} \in S$. There is some node $\omega \in S$ which is relevant for σ_α (otherwise σ_α is equivalent to σ_α^*). Since α cannot observe his rival's effort or output, and the probabilities of production have full support for every effort level, there exist $\omega \in S$ which is reached with positive probability under $(\sigma_\alpha, \sigma_\beta^*)$

Go from σ_α to σ_α^* through a sequence of transitions $\sigma_\alpha \equiv \sigma_\alpha^1 \rightarrow \dots \rightarrow \sigma_\alpha^{L+1} \equiv \sigma_\alpha^*$ as in Section 17.2.3. By Lemma 6 we have $Prob_\alpha^n(\sigma_\alpha^{l+1}, \sigma_\beta) > Prob_\alpha^n(\sigma_\alpha^l, \sigma_\beta)$ in the transition at ω ; and by (13) and (2), we have “ \geq ” everywhere else. This proves $Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta) > Prob_\alpha^n(\sigma_\alpha, \sigma_\beta)$.

Then the Theorem follows from (8).

17.5 Proof of Theorem 3

We shall prove Theorem 3 in the context of a wider class of games than $\Gamma(\Psi, B, n)$, which may be of some interest in their own right.

Consider a game Γ with players 1 and 2, who have finite pure-strategy sets denoted, w.l.o.g., by $S_1 \equiv \{1, 2, \dots, M\}$ and $S_2 \equiv \{1, 2, \dots, N\}$; and payoffs $h_\alpha : S_1 \times S_2 \rightarrow R$ for $\alpha = 1, 2$.

Payoff hypothesis:

There exists a constant K and (for $\alpha = 1, 2$) functions $g_\alpha^+ : S_1 \times S_2 \rightarrow R$ and $g_\alpha^- : S_\alpha \rightarrow R$ such that

$$\begin{aligned} (i) \quad & h_1(m, n) = g_1^+(m, n) - g_1^-(m) \\ (ii) \quad & h_2(m, n) = g_2^+(m, n) - g_2^-(n) \\ (iii) \quad & g_1^+(m, n) + g_2^+(m, n) = K \end{aligned}$$

for all $(m, n) \in S_1 \times S_2$.

Let Δ_1 and Δ_2 denote the mixed-strategy sets of 1 and 2, i.e.,

$$\Delta_1 = \{(\eta_1, \dots, \eta_M) : 0 \leq \eta_i \leq 1 \text{ and } \sum_{i=1}^M \eta_i = 1\}$$

$$\Delta_2 = \{(\xi_1, \dots, \xi_N) : 0 \leq \xi_i \leq 1 \text{ and } \sum_{i=1}^N \xi_i = 1\}.$$

Then the mixed extension $\tilde{\Gamma}$ of the pure strategy game Γ is defined via payoffs

$$h_\alpha(\eta, \xi) = \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j h_\alpha(i, j).$$

for all $(\eta, \xi) \in \Delta_1 \times \Delta_2$.

Lemma 7

Suppose Γ satisfies the Payoff Hypothesis. Then

$$(m, n) \text{ is a strict SE of } \Gamma \Rightarrow (m, n) \text{ is the unique SE of } \tilde{\Gamma}.$$

Proof

Suppose, to the contrary, that $(\eta, \xi) \in \Delta_1 \times \Delta_2$ is an SE of $\tilde{\Gamma}$ different from (m, n) . Since (m, n) is a strict SE of Γ , it follows that $h_1(m, n) > h_1(\eta, n)$ i.e.,

$$g_1^+(m, n) - g_1^-(m) > \sum_{i=1}^M \eta_i g_1^+(i, n) - \sum_{i=1}^M \eta_i g_1^-(i) \quad (14)$$

Similarly, from $h_2(m, n) > h_2(m, \xi)$ we get

$$g_2^+(m, n) - g_2^-(n) > \sum_{j=1}^N \xi_j g_2^+(m, j) - \sum_{j=1}^N \xi_j g_2^-(j),$$

which (using Payoff Hypothesis (iii), and the fact that $\sum_{j=1}^N \xi_j = 1$) may be rewritten

$$K - g_1^+(m, n) - g_2^-(n) > K - \sum_{j=1}^N \xi_j g_1^+(m, j) - \sum_{j=1}^N \xi_j g_2^-(j). \quad (15)$$

Now, because (η, ξ) is an SE of $\tilde{\Gamma}$ we have $h_1(\eta, \xi) \geq h_1(m, \xi)$, i.e.,

$$\sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) - \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^-(i) \geq \sum_{j=1}^N \xi_j g_1^+(m, j) - \sum_{j=1}^N \xi_j g_1^-(m)$$

i.e.,

$$\sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) - \sum_{i=1}^M \eta_i g_1^-(i) \geq \sum_{j=1}^N \xi_j g_1^+(m, j) - g_1^-(m). \quad (16)$$

Similarly from $h_2(\eta, \xi) \geq h_2(\eta, n)$ we get

$$\sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_2^+(i, j) - \sum_{j=1}^N \xi_j g_2^-(j) \geq \sum_{i=1}^M \eta_i g_2^+(i, n) - g_2^-(n).$$

which (again using Payoff Hypothesis (iii)) may be rewritten

$$K - \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) - \sum_{j=1}^N \xi_j g_2^-(j) \geq K - \sum_{i=1}^M \eta_i g_1^+(i, n) - g_2^-(n). \quad (17)$$

Rearranging equations (14) and (16), we have the following two inequalities

$$\begin{aligned} g_1^+(m, n) - \sum_{i=1}^M \eta_i g_1^+(i, n) &> g_1^-(m) - \sum_{i=1}^M \eta_i g_1^-(i) \\ g_1^-(m) - \sum_{i=1}^M \eta_i g_1^-(i) &\geq \sum_{j=1}^N \xi_j g_1^+(m, j) - \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j). \end{aligned}$$

Using the above two inequalities, we have

$$g_1^+(m, n) - \sum_{i=1}^M \eta_i g_1^+(i, n) > \sum_{j=1}^N \xi_j g_1^+(m, j) - \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j). \quad (18)$$

Rearranging equations (15) and (17), we have again the following two inequalities

$$\begin{aligned} \sum_{j=1}^N \xi_j g_1^+(m, j) - g_1^+(m, n) &> g_2^-(n) - \sum_{j=1}^N \xi_j g_2^-(j) \\ g_2^-(n) - \sum_{j=1}^N \xi_j g_2^-(j) &\geq \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) - \sum_{i=1}^M \eta_i g_1^+(i, n); \end{aligned}$$

and these two inequalities imply

$$\sum_{j=1}^N \xi_j g_1^+(m, j) - g_1^+(m, n) > \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) - \sum_{i=1}^M \eta_i g_1^+(i, n). \quad (19)$$

Rearranging (19), we have

$$\sum_{j=1}^N \xi_j g_1^+(m, j) - \sum_{i=1}^M \sum_{j=1}^N \eta_i \xi_j g_1^+(i, j) > g_1^+(m, n) - \sum_{i=1}^M \eta_i g_1^+(i, n)$$

which contradicts (18), proving Lemma 5.

Remark 4 The proof of lemma 2 also shows that, if the payoff hypothesis holds, SE have the “interchange property”: when (σ_1, σ_2) and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ are SE, so are $(\sigma_1, \tilde{\sigma}_2)$ and $(\tilde{\sigma}_1, \sigma_2)$. However these SE are *not* necessarily payoff-equivalent.

17.5.1 Completion of the Proof of Theorem 3

When α is ignorant of his rival β in Ψ , it is easy to check that α 's expected disutility $\sum_{\omega \in \Omega(T+1)} p^{(\sigma_\alpha, \sigma_\beta)}(\omega) d_\alpha((e_\alpha^\omega(\tau))_{\tau=1}^T)$ is independent of σ_β . (Here $\omega \equiv (e_\alpha^\omega(\tau), e_\beta^\omega(\tau), q_\alpha^\omega(\tau), q_\beta^\omega(\tau))_{\tau=1}^T$.)

Also notice that any affine transformation of von-Neumann Morgenstern utilities leaves invariant the set of mixed-strategy SE of $\tilde{\Gamma}$. W.l.o.g., we can therefore assume that $u_1(B) = u_2(B) = K$, a constant. Then the game $\Gamma(\Psi, B, n)$ satisfies the Payoff Hypothesis (since whenever 1 gets the bonus B with probability p , 2 gets it with probability $1 - p$). Theorem 3 now follows from Lemma 7.

17.6 Proof of Theorem 5

Denote $\mu_\alpha \equiv \sum_{q \in Q_\alpha} p_\alpha^{e_\alpha^*}(q)q$ and $\nu_\alpha \equiv \sum_{q \in Q_\alpha} p_\alpha^{e_\alpha^*}(q)[q - \mu_\alpha]^2$. Suppose that there is a sequence, indexed by k , such that $n(T_k) > cT_k$ for some $c > 0$. By Chebyshev's inequality, it follows that agent 2 wins the bonus with probability at most (noting $\mu_1 > \mu_2$ by Assumption VIII):

$$P\left[\left|\sum_{\tau=1}^{n(T_k)} (q_1(\tau) - q_2(\tau)) - n(T_k)(\mu_1 - \mu_2)\right| \geq n(T_k)(\mu_1 - \mu_2)\right] \leq \frac{n(T_k)(\nu_1 + \nu_2)}{(n(T_k))^2(\mu_1 - \mu_2)^2}$$

$$= M/n(T_k) \text{ where } M = (\nu_1 + \nu_2)/(\mu_1 - \mu_2)^2.$$

For σ_2^* to be the best reply to σ_1^* , B must satisfy: $u_2^{T_k}(B) \text{Prob}_2(\sigma_1^*, \sigma_2^*) > d_2^{T_k}(\bar{e}_2^*(T_k)) > T_k/\gamma$ (by (i) of Assumption VII), i.e., $u_2^{T_k}(B) > (T_k/\gamma)(\text{Prob}_2(\sigma_1^*, \sigma_2^*))^{-1} > (T_k/\gamma)(n(T_k)/M)$, i.e.,

$$u_2^{T_k}(B) > \frac{c}{M\gamma} T_k^2 \tag{20}$$

Let $p_\alpha(e_\alpha, e_\beta)$ denote the probability that α wins the bonus if he puts in effort e_α and β puts in e_β (when both work for 1 day). Denote

$$\delta \equiv \min_{\alpha \in \{1,2\}, e_\alpha \in \Sigma_\alpha \setminus \{e_\alpha^*\}} p_\alpha(e_\alpha^*, e_\beta^*) - p_\alpha(e_\alpha, e_\beta^*)$$

By assumption (II), $\delta > 0$. We shall show that if

$$u_\alpha^T(B) > \frac{\gamma}{\delta} T \quad (21)$$

then σ_α^* is the best reply to σ_β^* for $\alpha \in \{1,2\}$, when the sample size is 1. Consider the sequence $\sigma_\alpha \equiv \sigma_\alpha^1 \rightarrow \sigma_\alpha^2 \rightarrow \dots \rightarrow \sigma_\alpha^{l+1} \equiv \sigma_\alpha^*$ as in the proof of Lemma 2, with $\sigma_\alpha^{l+1} \succ_{\omega_l} \sigma_\alpha^l$. Let $p(\sigma_\alpha, \sigma_\beta, \omega)$ be the probability that node ω is reached under $(\sigma_\alpha, \sigma_\beta)$ and note $p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l) = p(\sigma_\alpha^{l+1}, \sigma_\beta^*, \omega_l)$. When α switches from σ_α^l to σ_α^{l+1} , his disutility goes up by at most $p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l)\gamma$ by (ii) of Assumption VII. On the other hand ω_l is sampled with probability $p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l)/T$ (since each day is sampled with probability $1/T$ when the sample size is 1). Thus the gain in utility to α , when he switches from σ_α^l to σ_α^{l+1} is at least $(p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l)/T)\delta u_\alpha^T(B)$. Hence (21) implies that α profits by the switch whenever $p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l) > 0$ (and is unaffected if $p(\sigma_\alpha^l, \sigma_\beta^*, \omega_l) = 0$). Then σ_α^* is the unique best reply to σ_β^* (since some ω_l are reached with positive probability). For large enough T , there exist $B(T)$ for which (21) holds while (20) is violated. (We use Assumption III, applied to u_α^T , here.) This proves that sample size 1 is better for the principal than sample size $n(T)$ for all large enough T , contradicting that $n(T)$ is an optimal size.

17.7 Proof of Remark 3

By (9), $n \leq g(\Psi)$ implies $Prob_\alpha^n(\sigma_\alpha^*, \sigma_\beta^*) > Prob_\alpha^n(\sigma_\alpha, \sigma_\beta^*)$ for $\sigma_\alpha \in \Sigma_\alpha \setminus \{\sigma_\alpha^*\}$. We claim that there is no $(\sigma_1, \sigma_2) \neq (\sigma_1^*, \sigma_2^*)$, such that (a) $Prob_1^n(\sigma_1, \sigma_2) \geq Prob_1^n(\sigma_1^*, \sigma_2)$ and (b) $Prob_2^n(\sigma_1, \sigma_2) \geq Prob_2^n(\sigma_1, \sigma_2^*)$. Otherwise, since $(\sigma_1, \sigma_2) \neq (\sigma_1^*, \sigma_2^*)$, assume w.l.o.g. $\sigma_1 \neq \sigma_1^*$. We have $Prob_1^n(\sigma_1^*, \sigma_2^*) > Prob_1^n(\sigma_1, \sigma_2^*) \geq Prob_1^n(\sigma_1, \sigma_2) \geq Prob_1^n(\sigma_1^*, \sigma_2) \geq Prob_1^n(\sigma_1^*, \sigma_2^*)$. (The strict inequality is from (9) and the fact that $\sigma_1 \neq \sigma_1^*$, the second inequality is from (b), the third inequality is from (a) and the last inequality is from Lemma (2).) This is a contradiction. Thus for every pair $(\sigma_1, \sigma_2) \neq (\sigma_1^*, \sigma_2^*)$, there exists $B(\sigma_1, \sigma_2)$ such that, if $B > B(\sigma_1, \sigma_2)$, at least one agent α would like to unilaterally deviate from σ_α to σ_α^* . Let $\hat{B} = \max_{(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2 \setminus \{(\sigma_1^*, \sigma_2^*)\}} \{B(\sigma_1, \sigma_2)\}$. Then for any $B > \max\{\hat{B}, B(\Psi, n)\}$, (σ_1^*, σ_2^*) is the unique SE.

References

- Bulow, J., Geanakoplos, J., and Klemperer, P. (1985). "Multimarket Oligopoly: Strategic Substitutes and Complements" *J. Polit. Econ.* **93**, 488-511.
- Cowen, T., and Glazer, A. (1996). "More Monitoring can Induce Less Effort" *J. of Economic Behavior & Org.* **30**, 113-123.
- Dubey, P., and Haimanko, O. (2000). "Optimal Scrutiny in Multi-Period Promotion Tournaments", Cowles Foundation Discussion Paper 1254, Yale University.
- Green, J. R., and Stokey, N. L. (1983). "A Comparison of Tournaments and Contracts." *J. Polit. Econ.* **91**, 349-64.
- Grossman, S. J., and Hart, O., (1983). "An Analysis of The Principal-Agent Problem." *Econometrica* **51**, 7-45.
- Lazear, E. P., and Rosen, S., (1981). "Rank-Order Tournaments as Optimum Labor Contracts." *J. Polit. Econ.* **89**, 841-64.
- Mookherjee, D., (1984). "Optimal Incentive Schemes with Many Agents." *Review of Economic Studies* **LI**, 433-446.
- Nalebuff, B. J., and Stiglitz, J. E., (1983). "Prizes and Incentives: Towards a General Theory of Compensation and Competition." *Bell J. Econ.* **14**, 21-43.
- Rosen, S., (1986). "Prizes and Incentives in Elimination Tournaments." *A.E.R.* **76**, 701-15.