Review of Economic Studies (1991) **58**, 15-41 0034-6527/91/00020015$02.00

Learning from Coarse Information: Biased Contests and Career Profiles

MARGARET A. MEYER
Nuffield College, Oxford

First version received January 1990; *final version accepted October* 1990 (*Eds.*)

An organization's promotion decision between two workers is modelled as a problem of boundedly-rational learning about ability. The decision-maker can bias noisy rank-order contests sequentially, thereby changing the information they convey. The optimal final-period bias favours the "leader", reinforcing his likely ability advantage. When optimally biased rank-order information is a sufficient statistic for cardinal information, the leader is favoured in every period. In other environments, bias in early periods may (i) favour the early loser, (ii) be optimal even when the workers are equally rated, and (iii) reduce the favoured worker's promotion chances.

1. INTRODUCTION

This paper presents a model of boundedly-rational learning by an organization. The aims are to represent in a simple form the costs of gathering, processing, or transmitting information and to analyse their implications for the procedures selected to handle information, as well as for the decisions actually made. Like other recent work on the internal organization of firms, this work explores to what extent organizational structure and behaviour, as well as organizational performance, can be explained by adaptation to the costs associated with information.[1]

In the learning model we formulate, the decision-maker not only selects a rule specifying what action to take as a function of his observations but also chooses, sequentially, the information partitions generating those observations. This general problem is interpreted in the context of an organization which needs to make an important promotion decision and which attempts to maximize the value of the information generated during the observation periods that precede the decision. One of the motivations for this analysis was the finding by organizational sociologists that earnings and promotions in the later stages of a worker's career are strongly correlated with earnings and promotions in the early stages (Kanter (1977), Rosenbaum (1984)). Later success is positively associated with early success, even when one controls for the effect of observable characteristics likely to affect performance, such as education. One obvious source of this correlation is differences in ability that persist over time and that are not captured by the observable covariates. In this paper, we focus on differences in ability which are initially unobservable by organizations and their workers, as well as by researchers. We present a simple model of an organization trying to learn about differences in ability and show

1. Influential discussions of these issues are in Arrow (1974) and Simon (1976). Recent analyses include Sah and Stiglitz (1986) and Geanakoplos and Milgrom (1988).

that the optimal learning strategy can reinforce the effects of these differences: early success can be even more strongly associated with later success under the optimal learning strategy than under a naive process of information accumulation.

Other explanations for the observed correlations, also based on an organization's response to limitations on its information about workers, can be developed. Meyer (1986) assumes that the effort levels of risk-averse workers are private information and shows that an organization can benefit by designing promotion schemes that reward early success with an increased probability of success in later contests. Although the consequent asymmetry in later contests between those with good and bad early records reduces incentives at that point, this effect is outweighed by the increase in incentives in early contests due to the future rewards for winning. Thus, even when abilities are known and equal, a correlation between early success and later success may be induced by the organization in an attempt to limit the costs of moral hazard.

In the learning model studied here, the effort decisions of workers are suppressed. Furthermore, workers are assumed to have no private information about their abilities. Because workers do not make any strategic decisions, the model is essentially a single-person decision problem, with the organization as the decision-maker. The organization's sole objective is to use the information generated during the fixed number of observation periods to identify and promote the more able of two workers.

We view the limitations on the information that becomes available to the organization as limitations on its "rationality", and we will analyse how the organization optimally responds to these limitations. We distinguish between the individual who designs the rules according to which the organization operates (the "decision-maker") and the individuals who implement the rules on a day-to-day basis (the "observers"). The decision-maker must be quite sophisticated to solve the problem of optimal organization design subject to the informational constraints. The observers embody these constraints in their limited observational ability or their limited ability to communicate what they observe.

In each period, an observer reports a signal from a very coarse partition of worker outputs: specifically, he reports only which of the two workers produced the larger output. In addition, output itself is affected by exogenous noise as well as by ability. Refining the coarse partition or reducing the exogenous noise is likely to be relatively costly for the organization. We focus instead on a means of changing the information structure which is relatively inexpensive and which can be accomplished through numerous commonly observed practices: adjusting the rule that determines whether the observer reports a "win" for one worker or the other. Instead of worker i being declared the winner whenever i's output exceeds j's output, i is declared the winner as long as his output does not fall short of j's by more than some amount c, and c is freely chosen by the decision-maker. We refer to c as the level of bias in i's favour.

Bias can be implemented in two different ways, depending upon the explanation for the coarseness of the observers' reports. When observers obtain cardinal information about outputs but, because of the costs of communication, report only ordinal information, the decision-maker can instruct observers to change the critical cutoff that determines whether they report a win by i or j, i.e. to use asymmetric evaluation criteria. On the other hand, when observers can identify only which of two output levels is larger, the level of bias can be controlled by the decision-maker by differentiating workers' tasks or work environments or by providing different amounts of training or equipment. Examples of the latter form of bias include the assignment of workers to clients with different needs or attitudes and the provision by senior colleagues of different levels of guidance.

Whatever the method of implementing bias, changes in its level change the information content of the observers' reports and hence affect the quality of the promotion decision. The problem facing the decision-maker is to choose the level of bias sequentially during the observation stage, along with a promotion rule specifying whom to select as a function of the biased observations.[2] Since the levels of bias will affect the probabilities with which workers win and lose each contest, the use of bias will have implications for the career profiles of workers within the organization.

Overview of the Paper

Section 2 uses a simple example to explain how biasing coarse and noisy observations can increase the value of the information they provide. Section 3 describes the model of learning in organizations. Section 4 shows that it is always optimal for the organization to bias the final contest before the promotion decision in favour of the current "leader", the worker with the better cumulative performance record. Thus, with respect to promotion chances, the advantage the leader derives from his likely edge in ability is reinforced by the organization's optimal learning strategy. The optimal bias in favour of the leader equals the smallest margin of victory by his rival which (if actually observed) would outweigh the leader's better record: in the final period, therefore, the bias allows the decision-maker to choose between the workers exactly as he would if he had cardinal information on the difference in outputs.

Section 5 studies the use of bias in early periods. We present a necessary condition for optimality (Section 5.1) and contrast it to the characterization of the optimal final-period bias. In early periods, optimally biased rank-order information is not in general as valuable as cardinal information, because the former is not in general a sufficient statistic for the latter with respect to decisions about future levels of bias. In Section 5.2, we characterize the stochastic environments in which the sufficient statistic property does hold and show that in these environments, the optimal strategy for setting bias is myopic: in each period, the bias should be set as if that period were the last, so success in the current contest is always reinforced through bias in the next one.

In stochastic environments in which the sufficient statistic property does not hold, bias may be used differently in early periods than in the final period. Section 6.1 presents an example in which an organization with three periods before the promotion benefits by rewarding in the second period the worker who performed *worse* in the first period. After explaining this possibility, we note that the desirability of favouring the leader in early contests would be increased if the model were changed to include not only the "major" promotion decision but also "minor" job assignment decisions during the observation phase: by increasing the importance, at each stage, of identifying the better worker, this change would make each period's problem more closely resemble the final-period problem.

Section 6.2 shows that an organization may improve its promotion decision by using bias in the first contest, before any differences in ability have been revealed. While notions of fairness would dictate the random assignment of first-period bias, there would be no loss in efficiency from assigning the bias according to economically irrelevant, but socially meaningful, criteria such as race, sex, or family background. Treating the workers asymmetrically in the initial contest is more likely to help the organization, the more

2. In the terminology of statistical decision theory (see, e.g. Raiffa and Schlaifer (1961)), the promotion choice is a "terminal act" and the biased contests are "experiments" which the decision-maker can adaptively design, according to previous observations.

sensitive are future biases to changes in first-period bias and the less costly it would be to introduce bias in a one-period choice problem.

Section 6.3 demonstrates that the introduction of bias in a worker's favour in an early period may reduce his overall chance of promotion, even when it benefits the organization. Introducing bias in a worker's favour lowers the decision-maker's degree of confidence in him whether he wins or loses the current contest. If this change in beliefs alters future levels of bias in a sufficiently adverse way, this future handicap can outweigh, in terms of its effect on promotion chances, the benefits from the increased likelihood of winning the current contest.

Section 7 considers the use of bias when, after the observation phase, the decision-maker is not restricted to the binary choice of which worker to promote. This extension introduces the same possibilities in the final period as arise in the original model in early periods, since later choices about levels of bias are not restricted to be binary. In an example in which the decision-maker has the option of making no promotion, we show that even in the final period, it can be optimal to treat symmetrically rated workers asymmetrically and, when there is a leader, to set the bias against him. When the decision-maker can choose only which worker to promote, he maximizes his expected degree of coinfidence in the promoted worker; adding the option of promoting no one if he is not sufficiently confident makes him prefer strategies which produce a very variable degree of confidence, even at the expense of some reduction in the mean.

Related Work

Statisticians have studied sequential sampling rules in problems such as whether to accept or reject a consignment of goods or which of two medical treatments to adopt (Kulkarni (1982), Armitage (1985)). In clinical trials, *adaptive* sequential procedures allow not only the stopping decision but also the choice of which treatment to use at any stage to be a function of past observations.[3] Evaluations of procedures consider not only the probability of selecting the better treatment but also the length of the trial and the number of volunteers given the inferior treatment. One formulation of the problem is as a two-armed bandit problem (Berry and Fristedt (1985)). In economics, related models have been used to analyse a firm trying to learn about its demand curve (Rothschild (1974), Aghion *et al.* (1990), Alpern and Snower (1990), and Balvers and Cosimano (1990)).

Two features differentiate our analysis from those above. First and most important, we stress that when observations are coarse (win/lose, succeed/fail), it is valuable for the decision-maker to adapt to past observations by altering his information partitions over outcomes (through bias), thereby altering the information content of the different coarse reports. In clinical trials, the analogue would be changing the criteria distinguishing success from failure, but such a possibility seems not to have been studied.

Second, we suppress both the direct and the opportunity costs of learning; bias is costlessly adjustable and has no effect on total output during the observation stage. Hence, in contrast to the analyses above, there is no tradeoff between short-run payoff maximization and the generation of information valuable for the future.[4]

3. For example, the "play the winner" rule (Zelen (1969)) tests the treatment used last period if it produced a success and the other treatment otherwise.

4. The opportunity cost of learning is crucial to the four economics papers on learning cited above, which investigate whether the firm's price converges to the full-information optimum and how the learning motive distorts decisions away from the myopic profit-maximizing ones.

Several recent models of decision-making in economics and political science share with ours the feature that information partitions or information sources are themselves choice variables.

Dow (1991) studies a boundedly-rational consumer who visits two stores in succession, searching for a low price. When deciding where to buy, the consumer cannot perfectly recall the first store's price, but only in which element of an n-element partition of the set of possible prices it lay. Before starting his search, the consumer anticipates this memory limitation and chooses the partition optimally. This problem of the optimal design of limited memory is formally similar to the choice of bias in early periods in our model.

Sah and Stiglitz (1986) analyse an organization deciding whether to undertake a project of uncertain value. Each of two evaluators observes a noisy signal of the value but reports only whether his signal exceeds a specified standard, which is freely adjustable by the organization. Sah and Stiglitz assume that the passing standards must be chosen *simultaneously* and focus on the resulting choice between requiring only one "pass" for acceptance (polyarchy) or two (hierarchy). Since we allow *sequential* adjustment of biases, neither the hierarchical nor the polyarchical decision rule is optimal in the two-period version of our problem; a worker will be promoted if and only if he wins the final contest.

Fishman and Hagerty (1990) study how much discretion an entrepreneur should be allowed in choosing what information to report to potential investors. At a formal level, changes in the permitted amount of discretion alter investors' information partitions in a manner very similar to bias in our model. The finding that the more likely a project is to be good, the more discretion the entrepreneur should be given parallels our final-period result that the greater the degree of confidence in the leader, the larger should be the bias in his favour.

Calvert (1985) analyses how a rational political decision-maker will choose among advisors who provide imperfect, binary recommendations about two policy options. When the decision-maker is predisposed towards one option, it is optimal to seek advice from a source who is biased in favour of that option: only with an advisor biased in this way can a negative recommendation about the initially preferred option be sufficiently convincing to induce the decision-maker to choose the other one. Thus a preference for information likely to confirm one's beliefs, a phenomenon of great interest to experimental psychologists (see Wicklund and Brehm (1976)), may be rational. Though Calvert uses a different formalization of bias, his finding, too, parallels our result on the optimality of final-period bias in favour of the leader.

In contrast to our analysis of bias in a sequence of contests, none of these models focuses on how information-gathering strategies will vary over time.[5] We emphasize that the quality of decisions can be improved by sequentially adjusting information partitions according to previous observations. Such adjustments cause the *interpretations* of coarse reports to change over time, even if the *labels* they are given (win/lose, succeed/fail) remain the same.

2. THE VALUE OF BIAS: AN EXAMPLE

A decision-maker, D, will win one dollar if he bets correctly on the outcome of a spinner that is coloured red and green. He knows that either (case R) two-thirds of the spinner's area (240°) is red and the rest (120°) green, or (case G) two-thirds (240°) is green and one-third (120°) red. These two configurations are equally likely.

5. See, however, Calvert's reference to choosing which of two advisors to consult first (p. 551).

An observer, O, conducts a test spin, reporting the colour on which the needle lands. The information thus provided about the true configuration is noisy: the final position of the needle is random, so it is likely, but not certain, to land on the colour occupying the larger area. The information is also coarse: O reports only "red" or "green", nothing more about the position of the needle. Despite its noisiness and coarseness, O's report is valuable to D: without the report, the probability that D bets in accord with the true configuration is $\frac{1}{2}$, whereas with the report, this probability is $\frac{2}{3}$.

Does D's expected payoff increase if O carries out a second test spin, identical to the first, before D places his bet? No. If the spins produce different results, D is indifferent between betting on red and green. Therefore, one of D's optimal strategies is to ignore the result of the second spin and to bet on the colour that turns up on the first spin.

However, a simple alteration to the spinner by O will make O's report on the second spin valuable, even if the needle continues to come to rest randomly and O continues to report only the colour on which the needle lands. This alteration is to enlarge the sector covered by the colour reported on the first spin. (The spinner is returned to its original configuration before D places his bet.) Suppose, for example, that the outcome of the first spin is red. Before the second spin, O converts a green sector of c degrees to red, whatever the true configuration, and then spins the needle and reports the colour on which it lands. With this change, the unique optimal betting strategy for D is to bet on the colour which O reports on the second spin:

$$\frac{P\,(\text{true configuration is G}\,|\,\text{red on 1st, green on 2nd})}{P\,(\text{true configuration is R}\,|\,\text{red on 1st, green on 2nd})}$$

$$=\frac{P(\text{green on 2nd}\,|\,\text{true configuration is G})}{P(\text{green on 2nd}\,|\,\text{true configuration is R})}\,\frac{P(\text{true configuration is G}\,|\,\text{red on 1st})}{P(\text{true configuration is R}\,|\,\text{red on 1st})}$$

$$=\left(\frac{240-c}{120-c}\right)\frac{1/3}{2/3}>1 \quad \text{for } c>0.$$

Since D's optimal bet depends on the outcome of the second spin when O modifies the spinner, the information from the second spin has strictly positive expected value for D.

To determine the optimal value of c, calculate the probability (as a function of c) that D bets in accord with the true configuration, when he bets on the colour reported on the second spin:

$$P(\text{red on 2nd}\,|\,\text{true configuration is R, red on 1st})P(\text{true configuration is R}\,|\,\text{red on 1st})$$
$$+P(\text{green on 2nd}\,|\,\text{true configuration is G, red on 1st})P(\text{true configuration is G}\,|\,\text{red on 1st})$$

$$=\left(\frac{240+c}{360}\right)\left(\frac{2}{3}\right)+\left(\frac{240-c}{360}\right)\left(\frac{1}{3}\right) \quad \text{for } 0\leqq c\leqq 120.$$

This probability is linearly increasing in c, so the optimal c in the range [0°, 120°] is 120°. As c is increased beyond 120°, D's expected payoff falls, because the probability that red comes up on the second spin in state R stays at 1 (the red sector can get no larger than 360°), but the probability that green comes up in state G is reduced. Therefore, the optimal alteration to the spinner makes it completely red if the true configuration is R (and two-thirds red if the true configuration is G). With this change, if O reports green on the second spin, D knows with certainty that the state is G. D's probability of betting in accord with the true state is $\frac{7}{9}$.

The lesson of this example is that by changing the information partition that determines what O reports on the second spin, according to the outcome of the first spin, D can increase the expected value of O's reports. The change that benefits D increases the likelihood that the colour that turns up on the first spin turns up again on the second.

3. A MODEL OF LEARNING BY AN ORGANIZATION

3.1. *Assumptions*

An organization employs a pair of workers (labelled i and j) for T periods, after which it selects one of the workers for promotion to a different job. In each period t before the promotion decision, worker k's output, x_k^t, is the sum of his time-invariant ability, η_k, and a time-dependent noise term, ε_k^t:

$$x_k^t = \eta_k + \varepsilon_k^t, \qquad k = i, j, \qquad t = 1, \ldots, T.$$

We assume that the promoted worker's output in the new job is more sensitive to his ability than is his output in the original job, so the organization wants to promote the more able worker.

The number of periods before the promotion decision is fixed at T. No change in job assignment is possible after the promotion decision. It is irrelevant for our analysis whether the worker who is not promoted continues to perform the original job or leaves the organization.

With the technology above, the quality of the promotion decision is constrained by the fact that in each period output is a noisy indicator of ability. We assume that it is also constrained by the fact that in each period, only rank-order information about the workers' outputs is available to the decision-maker. However, the decision-maker can, at the start of each period, costlessly adjust the criterion that determines which worker is labelled the winner in that period. Formally, the worker declared the winner in period t will be i if $x_i^t + c^t > x_j^t$ and j if $x_i^t + c^t < x_j^t$, where c^t, the "bias' in period t, is a choice variable which can depend on the biased rank-order information available from all previous periods.[6]

Before discussing two interpretations of the bias c^t, we make assumptions on the probability distributions of workers' abilities, η_k, and the exogenous noise terms, ε_k^t. Since only rank-order information about outputs is available, it is sufficient to focus on the distributions of the difference in abilities, $\Delta\eta \equiv \eta_i - \eta_j$, and the difference in noise terms in each period, $\Delta\varepsilon^t \equiv \varepsilon_i^t - \varepsilon_j^t$. Define $\Delta x^t \equiv x_i^t - x_j^t = \Delta\eta + \Delta\varepsilon^t$.

The workers are assumed to be ex ante indistinguishable, so the prior distribution of $\Delta\eta$ is symmetric about 0. For simplicity, we assume that $\Delta\eta$ can take on only two (equally likely) values, N and $-N$. We refer to the state $\Delta\eta = N$ ($\Delta\eta = -N$) as the state in which i (j) is the better worker. With a symmetric two-point support for $\Delta\eta$, the organization's objective of maximizing the expectation of output after the promotion decision is equivalent to maximizing the probability of promoting the better worker. The extension to an arbitrary (symmetric) distribution for $\Delta\eta$ is briefly discussed in Section 4.

The (differences in) exogenous shocks $\Delta\varepsilon^t$ are distributed independently across time and, for each t, symmetrically about 0. (We place no restrictions on the correlation, for a given t, between ε_i^t and ε_j^t.) Let $\Delta\varepsilon^t$ have a density, denoted $g^t(\cdot)$, that is continuous and strictly positive everywhere on the interior of its support $[-E^t, E^t]$, where E^t may be infinite. We assume that $E^t > N$, for all t. (If, for some t, $E^t < N$, the organization's

6. Under our assumptions below, the event $x_i^t + c^t = x_j^t$ will have zero probability.

decision problem would be trivial: in this period the effect of the noise would be so small that an unbiased rank-order observation would conclusively reveal which worker was more able.) We define $H^t(a) \equiv P(\Delta\varepsilon^t \geqq a) = \int_a^{E^t} g^t(z)dz$.

Finally, we assume that a monotone likelihood-ratio condition (MLRC) holds for inferences about $\Delta\eta$ from realized values of Δx^t. Formally, we assume that

$$P(\Delta\eta = N \,|\, \Delta x^t = -c) \text{ is decreasing in } c,$$

which is equivalent to the condition that the likelihood ratio for the observation $\Delta x^t = -c$ in the two states $\Delta\eta = N$ and $\Delta\eta = -N$ is a decreasing function of c:

$$\frac{P(\Delta x^t = -c \,|\, \Delta\eta = N)}{P(\Delta x^t = -c \,|\, \Delta\eta = -N)} = \frac{g^t(-c-N)}{g^t(-c+N)} \text{ is decreasing in } c. \tag{1}$$

Condition (1) says that, in the hypothetical case in which the decision-maker observed that i and j "tied" in period t with bias c in i's favour ($x_i^t + c = x_j^t$), this news would be worse news about i's ability, the larger was c. Condition (1) implies that

$$\frac{H^t(-c-N)}{H^t(-c+N)} \text{ is decreasing in } c. \tag{2}$$

Condition (2) says that the news that i won in period t with bias c in i's favour becomes worse news about i as c increases.

3.2. *Interpretations of bias*

The choice of the bias c^t can be interpreted in two ways, depending upon the explanation for why the decision-maker receives only ordinal information.

The first interpretation applies when the observer learns the actual values of outputs but can *transmit* ordinal information to the decision-maker far more cheaply than cardinal information.[7] The decision-maker can implement a level of bias c^t by instructing the observer to report that j is the winner in period t if and only if j's output exceeds i's by at least c^t.

The second interpretation applies when ordinal information about outputs is far less costly for the observer to *gather* than cardinal information.[8] In this setting, the decision-maker can control the level of bias by providing different inputs to the two workers' production functions or by differentiating the production functions themselves. The observer continues to observe and report only whether i's output is larger or smaller than j's, but these output values are perturbations on the x's above, which reflect only ability and exogenous noise: the observer compares $x_i^t + v_i^t$ with $x_j^t + v_j^t$, where v_i^t and v_j^t are chosen by the decision-maker so that $v_i^t - v_j^t$ equals the desired c^t.

Asymmetric treatment of workers, producing bias of the second type, takes numerous forms in practice. Workers may be assigned different tasks, placed in different environments, given different amounts of training or supervision, or supplied with different amounts of capital. For example, assistant professors being evaluated on their research can be burdened with different levels of administrative responsibility, junior lawyers can be given different amounts of guidance, and secretaries can be assigned equipment of different vintages. Furthermore, asymmetries often arise because jobs are not identical,

7. The costs of communication in organizations are stressed by Arrow (1974) and Stiglitz (1987).

8. An observer may, for example, be responsible for monitoring several pairs of workers and, given the overall limit on his attention, it may not be possible for him to acquire finer information about each pair.

and it would be difficult or costly to make them so (or to correct for the differences, given only coarse performance evaluations). For example, the needs of clients of a service firm may differ, but having employees split their time between different clients may incur high set-up costs. Or collaboration with senior colleagues may be essential and the personalities or talents of these individuals may differ, but frequent rotation of employees may be costly to organize.[9]

As these examples indicate, at least three complicating factors are ignored in the simple formalization of bias above. First, the decision-maker may not have complete control over the extent to which he differentiates the workers' production processes; the level of bias may, for example, be a random variable, with adjustable mean, and its realized value may not be observable. Second, adjusting bias by providing training or capital, or by having workers share different tasks, can be costly. Third, the levels of training and capital may interact multiplicatively, rather than additively, with ability (the marginal product of training or capital may be higher for a high-ability than for a low-ability worker), so the allocation of these resources between workers may affect total output even before the promotion decision. However, none of these factors would fundamentally alter the role of bias as a means of changing the decision-maker's coarse information partition on workers' outputs. The insights derived below into how the use of bias improves the promotion decision would remain valid even if we relaxed the assumptions that the level of c^t is deterministic, is costlessly adjustable, and interacts additively with ability.

More generally, when asymmetries in the treatment of workers arise from several sources and serve several functions, our analysis can be interpreted as showing the benefits of such asymmetries with respect to one particular problem faced by the organization—the selection of workers for promotion.

These benefits suggest an explanation for the finding by the organizational sociologists Baron and Bielby (1986) of a tendency of organizations to "fragment" work to a significant extent through the proliferation of job titles, "making finer distinctions among work roles than are required simply on the basis of job content". These small differences in job titles may be the organization's way of formally identifying the use of bias. This hypothesis is consistent with Baron and Bielby's finding that the proliferation of job titles is more extreme in firms in which workers' skills are highly firm-specific. The more firm-specific are the skills workers develop, i) the easier it is for a firm to introduce bias without driving away those whom the bias disadvantages and ii) the greater the firm's initial uncertainty about relative abilities, so the more heavily it will weight the learning motive in the design of job structures and promotion policies.

4. OPTIMAL BIAS IN THE FINAL PERIOD: FAVOUR THE LEADER

This section analyses the decision-maker's optimal choice of bias in the final period, T, before the promotion decision.

Suppose, to begin, that $T=2$, that the distributions of the noise terms are identical in the two periods ($g^1(\cdot)=g^2(\cdot)$), and that no bias is used in the first period ($c^1=0$). Then by analogy with the spinner example in Section 2, if we use no bias in the second

9. Bias could also arise from different actions taken by the workers themselves, whether these actions were subject to the control of the firm or privately chosen by the workers. In Rosen's (1986) analysis of incentives in elimination tournaments, for example, in pairings involving heterogeneous contestants, the individual (likely to be) of higher ability exerts more effort than his opponent. Such differences in efforts would affect the information about abilities that was conveyed by rank-order observations on outputs.

period, the rank-order report from the second period has no value for the promotion decision: it is an optimal strategy to ignore the second-round result and promote the winner of the first round. Similarly, if we introduce bias in the second period in favour of the loser of the first round (say j), then the optimal promotion rule selects i, whatever the second-round result: even if j wins the second round, condition (2) implies that the posterior probability that i is the better worker exceeds one-half. Hence, in this case, too, the information from the second round has no value. In contrast, with bias in the second round in favour of the first-round winner, i, the unique optimal promotion rule selects the winner of the second round: if j wins the second round, (2) implies that j is more likely to be better. In this case, as long as j has positive probability of winning the second round, the expected value of the second-round report is strictly positive.[10] Thus, with identical distributions of exogenous noise in the two periods, the organization can benefit by changing the partition on second-period output pairs to increase the likelihood that the first-period winner is declared the second-period winner as well.

We now generalize this conclusion. We consider the optimal choice of bias in period T, given an arbitrary history R^{T-1} of biases and rank-order observations in periods $1,\dots,T-1$. Define the "leader" in period T as the worker the decision-maker would promote if forced to choose at the *beginning* of period T, and let $\alpha^T \geqq \frac{1}{2}$ be the posterior probability, given the history R^{T-1}, that the period-T leader, say i, is better. With c^T freely variable, we can solve for the optimal c^T by assuming that the decision-maker uses the strategy of promoting the winner of round T; even the strategy of promoting the period-T leader can be expressed in this form by setting an infinite period-T bias in favour of the leader. The probability $Q^T(c^T)$ that the more able worker is promoted, when bias c^T is used in favour of the leader, i, is

$$
\begin{aligned}
Q^T(c^T) &= P(i \text{ wins round } T \text{ with bias } c^T \mid \Delta\eta = N, R^{T-1}) P(\Delta\eta = N \mid R^{T-1}) \\
&\quad + P(j \text{ wins round } T \text{ against bias } c^T \mid \Delta\eta = -N, R^{T-1}) P(\Delta\eta = -N \mid R^{T-1}) \\
&= H^T(-c^T - N)\alpha^T + H^T(c^T - N)(1-\alpha^T).
\end{aligned}
$$

Taking the hypothesis that the leader is more able as the null hypothesis, maximizing $Q^T(c^T)$ is equivalent to minimizing the expected total cost of Type I and Type II errors (which under our assumptions are equally costly).

Since $g^T(\cdot)$ is assumed continuous on $(-E^T, E^T)$, $Q^T(c^T)$ is differentiable for $c^T \in (-(E^T - N),\ E^T - N)$. For c^T in this range, the first-order condition can be rearranged to

$$
\frac{g^T(-c^T - N)}{g^T(-c^T + N)} = \frac{1-\alpha^T}{\alpha^T}.
$$

The left-hand side, the likelihood ratio for the observation $\Delta x^T = -c^T$, equals 1 at $c^T = 0$, is continuous for $c^T \in (-(E^T - N),\ E^T - N)$, and by the MLRC (condition (1)) is decreasing in c^T. Hence, as long as $\alpha^T > \frac{1}{2}$, an optimal c^T must be strictly positive. If the MLRC is strict, the optimal c^T is unique and strictly increasing in α^T. (If uniqueness does not hold, the optimal c^T values are an interval.)

10. The decision-maker's payoff as a function of c^2 fails to be concave at $c^2 = 0$. This is not strictly an example of the Radner-Stiglitz (1984) nonconcavity in the value of information, however, since c^2 is not a measure of the informativeness of the second contest: the payoff function is typically decreasing for c^2 sufficiently large.

When we allow for the possibility of a corner solution ($c^T = E^T - N$), an optimal c^T, given the MLRC, is characterized by[11]

$$\frac{g^T(-a-N)}{g^T(-a+N)} \geqq \frac{1-\alpha^T}{\alpha^T} \quad \forall a < c^T \qquad \text{and} \qquad \frac{g^T(-a-N)}{g^T(-a+N)} \leqq \frac{1-\alpha^T}{\alpha^T} \quad \forall a > c^T. \tag{3}$$

To interpret the optimal level of bias, imagine that in the final period the decision-maker could base his promotion decision on the actual value of Δx^T. An optimal promotion strategy, which minimized the expected total cost of Type I and Type II errors, would, by the MLRC, be described by a cutoff value $-d^T$: if $\Delta x^T > -d^T$, select i and if $\Delta x^T < -d^T$, select j.[12] As long as $\alpha^T > \frac{1}{2}$, any optimal cutoff $-d^T$ would be strictly negative: even if i lost the final round, the decision-maker would still choose to promote i if the margin of j's victory were sufficiently small.[13] When the decision-maker receives only rank-order information in the final period but can adjust the bias, a small bias in favour of i must improve the promotion choice relative to zero bias, because it causes i, rather than j, to be promoted for small, negative values of Δx^T. Setting $c^T = d^T$ is optimal, because it ensures that i is promoted if $\Delta x^T > -d^T$ and j if $\Delta x^T < -d^T$, so the decision-maker makes the same decision as if he could observe Δx^T directly.

Thus, the optimal final-period bias in favour of the leader (when unique) is the smallest margin of victory by the other worker which, if actually observed by the decision-maker in period T, would outweigh the leader's superior record and induce the decision-maker to promote the other worker. The greater the weight of the evidence from periods $1, \ldots, T-1$ in favour of the leader, the larger this critical margin of victory and therefore the larger the optimal final-period bias in the leader's favour.

This discussion is summarized in:

Proposition 1. (a) *By using bias optimally in the final contest, the decision-maker can make the same promotion decision, for each realization of* Δx^T, *as if cardinal information on* Δx^T *were available.*

(b) *Any optimal value of final-period bias favours the final-period leader by a strictly positive amount.*

(c) *For* c^T *in the range where* $Q^T(c^T)$ *is differentiable, a necessary and sufficient condition for optimality is: If a tie occurred and were observed by the decision-maker* ($x_i^T + c^T = x_j^T$), *he would be indifferent as to which worker to promote.*

(d) *When the optimal final-period bias is unique and characterized by* (c), *it is strictly increasing in* α^T, *the probabiity that the final-period leader is more able.*

It is easy to show that (a), (b), and (c) (as well as Proposition 2) remain valid if the difference in workers' abilities, $\Delta\eta$, has an arbitrary, rather than a two-point, (symmetric) distribution: in this more general setting, though, the maximization of expected post-promotion output does not reduce to the maximization of the probability of promoting the better worker, so the indifference condition in (c) above must be interpreted in terms of expected output.

11. We would need to work with (3) even for $c^T \in (-(E^T - N), E^T - N)$ if we relaxed the assumption of continuity of $g^T(\cdot)$. In this case, the optimal bias when $\alpha^T > \frac{1}{2}$ would never be negative, but could be 0, if the likelihood ratio were discontinuous at $c^T = 0$ and α^T were sufficiently close to $\frac{1}{2}$.

12. If the MLRC were strict, the optimal $-d^T$ would be unique. If uniqueness did not obtain, the set of optimal cutoff values would be an interval.

13. *Strict* negativity of $-d^T$ follows from the continuity of the likelihood ratio, $g^T(\Delta x^T - N)/g^T(\Delta x^T + N)$, at $\Delta x^T = 0$.

Asymmetric treatment of competitors according to recent performance is a common feature of sports contests. In tennis tournaments, for example, players with strong records are favoured by being seeded, which prevents them from competing in early rounds against others with strong records. Furthermore, among the seeds, those seeded higher are more advantaged, because they are further in the draw from other highly seeded players. In motor races, more favourable starting positions are allocated to drivers who have performed well in preliminary heats. Organizers of sports contests may have numerous objectives in addition to identifying the ablest contestants (either in the short run or in the long run). Nevertheless, the analysis in this section suggests that, as long as the advantages given to recent strong performers are not too large, one of their effects is to enhance the likelihood that the current contest is won by the competitor of the highest ability.

5. CHARACTERIZATION OF OPTIMAL BIAS IN EARLY PERIODS

Section 4 established that an optimal learning strategy for the organization employs bias in the final period that increases the leader's likelihood of promotion. The optimal bias ensures that the organization makes the same decision, given only ordinal information on Δx^T, as if it had access to cardinal information on Δx^T. This section contrasts the considerations governing the use of bias in early periods with those in the final period. By deriving a necessary condition for optimality (Section 5.1), we show why the bias in an arbitrary period $t<T$ will typically differ from the level that would be chosen if t were the final period. We also explain why the sequentially optimal use of bias cannot typically render ordinal information as valuable to the decision-maker as cardinal information on Δx^t in all periods t. In Section 5.2, however, we characterize the stochastic environments in which optimally biased rank-order information *is* a sufficient statistic for cardinal information, and we show that in these environments, it *is* optimal to compute the bias in each period as if that period were the last.

5.1. *A necessary condition for optimal bias in period t*

In an arbitrary period $t<T$, denote the history of biases and rank-order observations by R^{t-1}. Given R^{t-1}, define the leader in period t as the worker the decision-maker would promote if forced to choose at the beginning of period t, and let α^t denote the probability that the period-t leader is better. The decision-maker's choice of c^t affects not only the manner in which Δx^t values are partitioned into the sets for which i and j are declared the winner of round t but also the future beliefs for each of the two possible period-t observations, hence the choices of bias in future periods $\tau>t$. This latter effect is absent in period T, when the choice of c^T affects only which worker is declared the winner and thereby promoted.

The following proposition uses dynamic programming to extend the optimality condition for final-round bias to an arbitrary period $t<T$.

Proposition 2. *Consider an arbitrary $t<T$ and assume that for $\tau=t+1,\ldots,T$, the bias c^τ will be chosen optimally, given the history $R^{\tau-1}$. For c^t in the range where the objective function is differentiable, a necessary condition for optimality is: If a tie occurred and were observed by the decision-maker ($x_i^t+c^t=x_j^t$), he would be indifferent as to which worker to*

declare the winner in period t, given how the values of c^τ, $\tau = t+1, \ldots, T$, *would depend on* c^t *and on who was declared the period-t winner.*[14]

Proof. See Appendix. ‖

The necessary condition for optimality of c^t in an early period $t < T$ is not identical to the final-period condition: the requirement in Proposition 2 that (if the decision-maker were to observe that $\Delta x^t = -c^t$) he be indifferent as to whom to declare the winner of round t is not equivalent to the requirement in (c) of Proposition 1 that, after period t, he assign equal likelihood to each of the two workers being better. These two requirements would be equivalent in the special case in which the optimal future treatment of the round-t winner was independent of whether he was i or j. In this case, the optimal c^t would maximize the probability that the round-t winner was better, so c^t would be chosen as if the promotion were to occur immediately after round t (see Section 5.2). In general, however, the decision-maker will hold different beliefs in period $t+1$ about the period-t winner, depending upon whether he was i or j, and with different beliefs, the decision-maker will in general prefer to set different biases in periods $t+1, \ldots, T$ (see Section 6.1). Therefore, the optimal level of bias in period $t < T$ will typically differ from the level that would be chosen by a myopic strategy, i.e. one which set the bias as if t were the final period.

To gain further insight into Proposition 2, recall that we identified the optimal bias in the final period by imagining that the decision-maker could base his promotion decision on the actual value of Δx^T. In this case, for some d^T, he would wish to promote i (j) if $\Delta x^T > -d^T (\Delta x^T < -d^T)$. Setting a bias $c^T = d^T$ was optimal, because it implemented the optimal cardinal-information strategy. In an arbitrary period $t < T$, the analogous procedure is to imagine that the decision-maker could base the choice of the round-t winner on the actual value of Δx^t but that his future actions could depend only on the period-t bias and rank-order result (as well as on the history R^{t-1}). Then by the MLRC, his preferences could again be described by a cutoff value, now denoted $-d^t(c^t)$: he would wish to declare i (j) the period-t winner if $\Delta x^t > -d^t(c^t)(\Delta x^t < -d^t(c^t))$. Crucially, an optimal cutoff value $-d^t(c^t)$ would now depend not only on R^{t-1} but also on c^t itself, since c^t would affect the ultimate promotion decision through its influence on future beliefs and therefore on future optimal choices of bias.[15] An optimal value of c^t must satisfy the fixed-point property $c^t = d^t(c^t)$. That is, in determining the winner of round t, c^t must partition the values of Δx^t exactly as the decision-maker would wish to partition them, if the choice of the period-t winner, but not future choices, could be based on cardinal information about Δx^t.

To confirm that $c^t = d^t(c^t)$ is necessary for optimality, note that if future biases are set optimally given c^t, then by the envelope theorem, the first-order effect on the decision-maker's objective function of a small change in c^t can be identified by treating d^t as unchanged. Therefore, if $c^t = d^t(c^t)$, a small change in c^t has no first-order effect, whereas if $c^t \neq d^t(c^t)$, then a first-order improvement results from a small adjustment of c^t in the direction of $d^t(c^t)$.

The first-order condition for c^t, stated in words in Proposition 2 and represented by $c^t = d^t(c^t)$, is necessary, but not sufficient, for optimality: because of the dependence of

14. If differentiability fails, an optimal c^t must be such that, if $\Delta x^t > -c^t$ ($\Delta x^t < -c^t$), the decision-maker would be willing to declare i (j) the winner.

15. In this discussion, we assume for simplicity that for each c^t, the optimal d^t is unique (as is true if the MLRC is strict).

the decision-maker's preferred cutoff value, $d^t(c^t)$, on the value of c^t itself, there may be multiple, isolated solutions to $c^t = d^t(c^t)$. A solution is a local maximizer (local minimizer) if a small change in c^t causes a smaller (larger) change in $d^t(c^t)$. To check this, note that from the new value of c^t, a first-order improvement can be achieved in the former (latter) case by moving c^t towards (away from) its original value. In the final period, the preferred cutoff d^T does not depend on c^T, so the first-order condition $c^T = d^T$ (stated in Proposition 1) is both necessary and sufficient for optimality.

In an arbitrary period $t < T$, the ability to adjust the bias freely allows the decision-maker to overcome the limitation to ordinal information on Δx^t *only* with respect to the determination of the period-t winner, not with respect to future choices about biases and promotion. If future choices of bias would in fact assume more than two values in some period if they could vary with the actual value of Δx^t, then optimally biased rank-order information in period t is not a sufficient statistic for, and so is strictly less valuable than, cardinal information on Δx^t. Proposition 4 in the next subsection characterizes such situations. In these cases, the decision-maker would benefit from the ability to choose a partition of Δx^t values into n, rather than just 2, intervals.[16] Proposition 2 generalizes to this case as follows: for a boundary between two intervals to be optimal, it must be that if the decision-maker actually observed a realization of Δx^t on the boundary, he would be indifferent as to which of the adjacent intervals to classify the outcome in, given how the partitions chosen in future periods would depend on how the current value of Δx^t was classified.[17]

When biases can be freely and sequentially adjusted, the value of a given contest outcome and its effect on beliefs depend on the history of biases and rank-order results, because this history influences the choice of bias in the current contest. One implication is that if the tasks performed by the workers vary from contest to contest, so contests involve different distributions of exogenous noise, then the quality of the optimal promotion decision may depend on the order in which the tasks are performed. Corollary 1 of Proposition 4 in the next subsection illustrates this possibility. (In contrast, if cardinal information on Δx^t is available in all periods, or if only unbiased rank-order information can be obtained, then the quality of the optimal decision will be independent of the order of the observations.)

A second implication of the dependence noted above is that, unlike the case where bias is unavailable, the decision-maker's degree of confidence in the worker he ultimately promotes is not necessarily higher if this worker was the winner of a given early contest than if he was the loser. For example, if $T = 2$ and bias is used optimally, a loss in the first contest followed by a win in the second may be a more favourable signal about ability than wins in both contests: those who overcome the obstacles placed in the way of slow starters may prove themselves more effectively than those who succeed at both stages. This possibility, too, is illustrated in Section 5.2.

5.2. *When are optimally biased contests as valuable as cardinal information?*

This subsection analyses the optimal use of bias in the special case in which in each period, $\Delta\varepsilon^t$ is uniformly distributed on $[-E^t, E^t]$.[18] In this case, the model takes the

16. The MLRC ensures that even if the decision-maker were allowed to choose arbitrary n-element partitions of Δx^t values, the optimal partitions would consist of intervals.

17. A similar interpretation can be given to the optimality condition derived by Dow (1990) for the partition representing limited memory.

18. This case arises when the shocks ε_i^t and ε_j^t to workers' outputs are uniformly distributed and prefectly negatively-correlated.

same mathematical form as the spinner example in Section 2. We show that for this case, it is optimal, in any period t, for any history R^{t-1}, to compute the bias as if period t were the final period, and we demonstrate that T periods of optimally biased rank-order observations are as valuable to the decision-maker as cardinal information on Δx^t in every period. We also show that the latter result holds if and only if the noise terms $\Delta\varepsilon^t$ are uniformly distributed in all periods after the first.

For the uniform distribution, the likelihood ratio for the observation Δx^t in the two states $\Delta\eta = N$ and $\Delta\eta = -N$, $g^t(\Delta x^t - N)/g^t(\Delta x^t + N)$, assumes only three values as Δx^t varies (see Figure 1). For $\Delta x^t \in I_2^t \equiv [-(E^t - N), E^t - N]$, the likelihood ratio equals 1, so observing $\Delta x^t \in I_2^t$ would leave beliefs about relative abilities unchanged. For $\Delta x^t \in I_1^t \equiv [-(E^t + N), -(E^t - N))$, the likelihood ratio equals 0: a value of $\Delta x^t \in I_1^t$ cannot arise if $\Delta\eta = N$, so such an observation would reveal conclusively that $\Delta\eta = -N$. Correspondingly, for $\Delta x^t \in I_3^t \equiv (E^t - N, E^t + N]$, the likelihood ratio equals ∞: observing $\Delta x^t \in I_3^t$ would reveal conclusively that $\Delta\eta = N$.

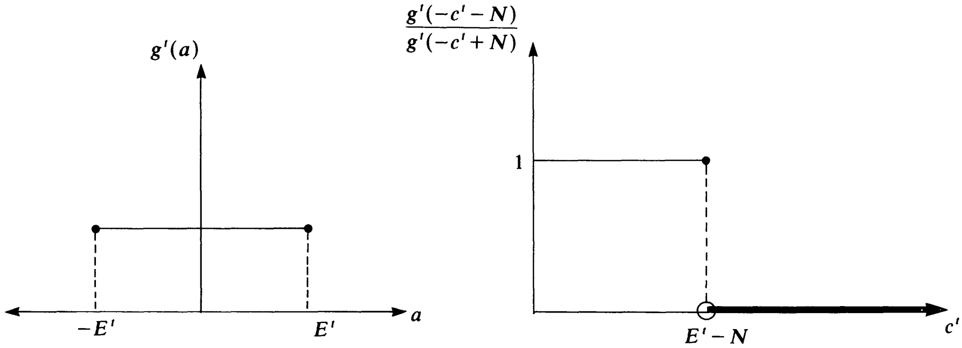

FIGURE 1
The uniform example

It follows, then, from condition (3) that in the final period, T, the bias should be set according to:

- If $\alpha^T \in (\frac{1}{2}, 1)$, favour the leader in period T by $E^T - N$.
- If $\alpha^T = 1$, favour the leader in period T by any amount greater than or equal to $E^T - N$.
- If $\alpha^T = \frac{1}{2}$, use any bias in the interval $[-(E^T - N), E^T - N]$.

Thus a bias of $E^T - N$ in favour of the leader is optimal, whatever the decision-maker's degree of confidence in the leader. Note that this level of bias makes the leader certain to win if he is in fact the more able worker.

Proposition 3 *Assume that for each $t = 1, \ldots, T$, the exogenous noise term $\Delta\varepsilon^t$ is uniformly distributed on $[-E^t, E^t]$.*

(a) *By using bias optimally, the decision-maker can achieve the same probability of promoting the better worker, for all realizations of $(\Delta x^1, \ldots, \Delta x^T)$, as if cardinal information on Δx^t were available in all periods.*

(b) *It is an optimal strategy to promote the winner of round T, to set the bias in period t, $t \geqq 2$, to favour the leader by $E^t - N$, and to choose any first-period bias in the interval $[-(E^1 - N), E^1 - N]$. With this strategy, the leader in period t is the winner of round $t-1$, and for all $t \geqq 2$ and all histories R^{t-1}, $\alpha^t > \frac{1}{2}$.*

(c) *Any strategy which sets the bias myopically, i.e. as if period t were the final period, is optimal.*

(d) *Suppose that when the workers compete at task k, $k = 1, \ldots, T$, the support of the exogenous noise is $[-E^k, E^k]$, regardless of when in the T periods this task is actually performed. Then the decision-maker's probability of promoting the better worker after T periods, when bias is used optimally, is independent of the order in which the tasks are performed.*

Proof. (a) and (b) Since the likelihood ratio for the observation Δx^t equals 0 if $\Delta x^t \in I_1^t$ and equals ∞ if $\Delta x^t \in I_3^t$, a necessary condition for an optimal promotion policy when cardinal information on Δx^t is available in all periods is

(A) j is promoted if for some t, $\Delta x^t \in I_1^t$, and i is promoted if for some t, $\Delta x^t \in I_3^t$. Since the likelihood ratio equals 1 if $\Delta x^t \notin (I_1^t \cup I_3^t)$, condition (A) is also sufficient for optimality, given that the decision-maker has symmetric prior beliefs about the workers.

The strategy described in (b) implements a promotion policy satisfying condition (A): if, for some t, $\Delta x^t \in I_1^t$ ($\Delta x^t \in I_3^t$), j (i) is certain to be declared the winner of round t and of all subsequent rounds, and is therefore certain to be promoted. Since the decision-maker cannot do strictly better by using biased rank-order observations than by optimally using cardinal information, the rules in (b) are an optimal strategy. This proves (a) and (b).

(c) Inspection of the optimal rules for final-period bias shows that any strategy satisfying them in each period implements a promotion policy satisfying condition (A).

(d) Since Part (a) holds whatever the order of the tasks and since the probability of promoting the better worker is independent of the order when cardinal information is used optimally, this probability is also independent of the order when bias is used optimally. ||

The analogue, in the spinner example of Section 2, of receiving cardinal information on Δx^t is learning the angular position of the needle (relative to a specified origin), as well as the colour on which it landed. Three regions can be identified on the spinner: A_R, which is known to be red; A_G, which is known to be green; and B, which is either red or green, depending on the true state. Because all angular positions are equally likely, whatever the state, learning that the needle landed in A_R or A_G is uninformative. However, learning that it landed in B and on what colour is conclusive. By Proposition 3, the optimal strategy for biasing the test spins when only the colour is reported is to convert all of A_G (A_R) to red (green) when state R (state G) is believed more likely; this strategy does as well as when both the angular position and the colour are reported.

For the uniform example, under an optimal strategy in which the bias is computed myopically, the outcome of the first contest establishes an initial leader and, for $t \geqq 2$, the bias partitions the values of Δx^t into those which, if actually observed by the decision-maker, would leave the identity of the leader unchanged and those which would conclusively establish that the leader was less able. The bias does not allow the decision-maker to distinguish between those values of Δx^t, for $t \geqq 2$, which reveal nothing about relative abilities and those which conclusively establish that the leader is more able. However, this additional information would never induce the decision-maker to alter his promotion choice and so is of no value for his decision.

With bias set optimally, the leader in any period after the first is certain to win if he is the more able worker, so if the leader loses, he is conclusively identified as less able.[19]

19. If the leader loses, there is no informational benefit from waiting until after period T to make the promotion.

Thus, the identity of the leader can change at most once during the T periods, and the only type of error the organization ever makes is to promote the less able worker when he wins the first and every subsequent contest (Type II error). The uniform example demonstrates that overcoming the bias which disadvantages slow starters can provide more convincing evidence of high ability than early and continued success.

In principle, the organization could undertake costly investments, in monitoring or communication technology, to refine the information available to the decision-maker about the difference in outputs, Δx^t. The following proposition generalizes (a) of Proposition 3 to characterize the environments in which such refinements would *not* improve the quality of the promotion decision.

Proposition 4. *If and only if the noise term $\Delta\varepsilon^t$ is uniformly distributed in each period after the first, the decision-maker can, by using bias optimally, achieve the same probability of promoting the better worker, for all realizations of $(\Delta x^1, \ldots, \Delta x^T)$, as if cardinal information on Δx^t were available in all periods.*

Proof. See Appendix. ||

The key to the proposition is the fact that the uniform distribution is the only distribution for which every value of Δx^t either reveals conclusively which worker is better or reveals nothing about relative abilities.

Proposition 4 provides a simple example of when a decision-maker who uses bias optimally can influence the quality of his promotion choice by adjusting the order in which the workers perform a given sequence of tasks:

Corollary 1. *If the noise terms are uniformly distributed for all except one task, k, then the ex ante probability of promoting the better worker is strictly higher when task k is performed first than when it is performed later in the sequence.*

Corollary 2. *If the noise term $\Delta\varepsilon^\tau$ is uniformly distributed for all $\tau > t$, then for any beliefs at the start of period t, it is optimal to set the period-t bias as if t were the final period.*

Proof. See Appendix. ||

When the noise terms in all future periods are uniformly distributed, the decision-maker will in future periods favour the leader by an amount that is independent of his degree of confidence. In this case, his objective in the current period reduces to maximizing the probability that the better worker wins the current contest (and so becomes next period's leader). This is the same objective as he would have if the current round were the final one.

Corollaries 3 and 4 analyse the benefits of improving the information about Δx^t in a single period, when optimally biased rank-order contests are used in all subsequent periods.

Corollary 3. *A refinement in the decision-maker's information partition on Δx^t can increase his ex ante probability of promoting the better worker only if, in some period τ, $\tau > t$, the noise term $\Delta\varepsilon^\tau$ is not uniformly distributed.*

Proof. See Appendix. ||

The converse of Corollary 3 is not valid.[20] However, a partial converse is given by:

Corollary 4. *Suppose that $T=2$ and that the bias in the second contest is adjusted optimally given the information from period 1. If $\Delta\varepsilon^2$ is not uniformly distributed, then the decision-maker's ex ante probability of promoting the better worker is strictly increased when he is allowed to partition the values of Δx^1 into three intervals of his own choosing, rather than just two.*

Proof. See Appendix. ||

6. FEATURES OF BIAS IN EARLY PERIODS

Section 4 showed that the optimal bias in the final period makes ordinal information as valuable as cardinal information on Δx^T and favours the leader. Section 5.2 proved that when T periods of biased rank-order observations are as valuable as T periods of cardinal information, then the bias should favour the leader in every contest after the first, and in the first, symmetric treatment of ex ante identical workers is optimal. We now demonstrate how the role of bias in early periods can differ when optimally biased rank-order information is not a sufficient statistic for cardinal information.

6.1. *Favouring the early loser may be beneficial*

Let $\Delta\varepsilon^t$ be distributed in each period as the difference of two independent and identically distributed exponential random variables, each having parameter λ^t (see Figure 2):

$$g^t(a)=\begin{cases}\dfrac{\lambda^t}{2}\exp(-\lambda^t a) & \text{if } a\geqq 0\\[2ex] \dfrac{\lambda^t}{2}\exp(\lambda^t a) & \text{if } a<0\end{cases}$$

$$H^t(a)=\begin{cases}\tfrac{1}{2}\exp(-\lambda^t a) & \text{if } a\geqq 0\\ 1-\tfrac{1}{2}\exp(\lambda^t a) & \text{if } a<0.\end{cases}$$

We show for a three-period problem that introducing a small bias in the second period in favour of the first-period loser may improve the promotion decision. Hence, in early periods, the organization may benefit (locally) from qualitatively different behaviour (rewarding early failure) than is optimal in the final period and, in the uniform case, optimal in every period (rewarding cumulative past success).

Lemma 1. *Assume that $T=3$, $\lambda^1=\lambda^2$, and i won the first round. Let $c^1=c^2=0$. If j is declared the winner of the second round, the optimal value of c^3, denoted c^3_{WL}, is 0. If i is declared the winner of the second round, the optimal value of c^3, denoted c^3_{WW}, is infinite (so i is promoted with certainty), as long as λ^3 is not too much larger than λ^1 (as long as $\lambda^3<(1/N)\ln(2\exp(\lambda^1 N)-1)\equiv\bar{\lambda}^3(\lambda^1)$, where $\bar{\lambda}^3(\lambda^1)>\lambda^1$).*

20. To construct a counterexample, let the noise terms be uniformly distributed in all periods except τ, for some $\tau>t$. If $\Delta\varepsilon^\tau$ has the exponential distribution of Section 6.1 and if biased contests are used in all periods other than t, it can be optimal to ignore the information from period τ, whatever the history of observations. Since with period τ ignored, the situation is equivalent to one in which the noise term is uniformly distributed in all periods after t, Corollary 3 implies that there is no benefit to refining the partition on Δx^t.

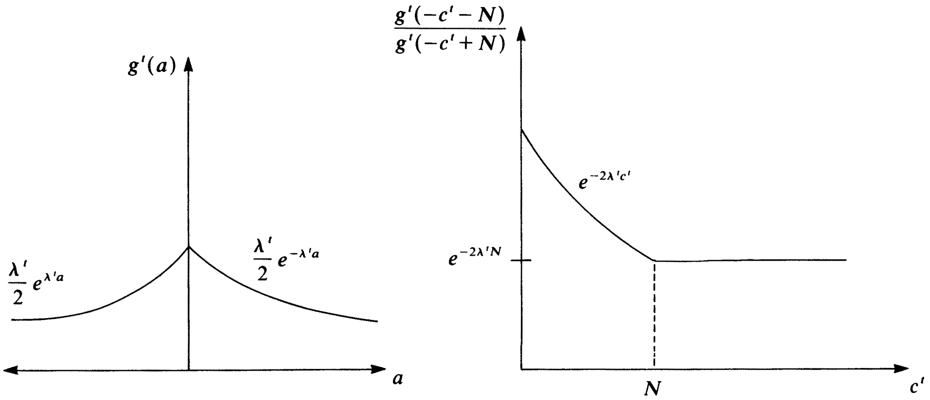

FIGURE 2
The exponential example

Proof. See Appendix. ||

Lemma 1 shows that the optimal treatment of the winner of the second contest is very sensitive to whether or not this worker also won the first contest.

Lemma 2. *Under the assumptions of Lemma* 1, *if* $\lambda^1 < \lambda^3 < \bar{\lambda}^3(\lambda^1)$, *the organization benefits from a small reduction in* c^2, *favouring the first-period loser j, starting from* $c^1 = c^2 = 0$. *If* $\lambda^3 < \lambda^1$, *the organization benefits from a small increase in* c^2, *favouring the first-period winner i, starting from* $c^1 = c^2 = 0$.

Proof. See Appendix. ||

Since $c^1 = 0$, $c^3_{WW} = \infty$, and $c^3_{WL} = 0$, Lemma 2 shows that the beneficial adjustment to c^2 increases the likelihood that the promotion is awarded to the winner of the unbiased contest in whichever of periods 1 and 3 is less noisy. Note that when a small bias in favour of the first-period loser is beneficial, this worker will not become the leader in period 3 even if he wins the second round; nevertheless, direct calculation shows that the bias increases his overall probability of promotion as well as his probability of winning the second round.

The possibility that the organization may benefit (locally) from disadvantaging the early leader arises because the size of the third-period bias is sensitive to the past performance of the second-round winner. If the optimal values of c^3_{WL} and c^3_{WW} satisfied $c^3_{WL} = -c^3_{WW}$, as they would if $\Delta\varepsilon^3$ were uniformly distributed, the decision-maker would prefer to declare as winner of the second round whichever worker was more likely to be better, so he would set the second-period bias as if $T = 2$: he would favour the first-period winner. The very large difference in this example between the magnitudes of $c^3_{WL} = 0$ and $c^3_{WW} = \infty$ is the reason why the decision-maker may prefer to disadvantage, rather than favour, the early leader.

The contrast between this example and the uniform case of Section 5.2 suggests that the less sensitive is the optimal future treatment of the winner of a given early contest to his previous history (when no bias is used in that contest), the more likely that the organization gains by introducing bias in that contest in favour of the current leader.

The desirability of favouring the leader in early contests would also be increased if the decision-maker faced, in addition to the "major" promotion decision after period T, "minor" job assignment decisions during the observation phase. These "minor" decisions would increase the importance, *at each stage*, of identifying the better worker, thereby making the decision-maker's problem at each stage more closely resemble his problem in the final period.

6.2. *Unequal treatment of equals may be optimal in an early period*

In the final period, a decision-maker with symmetric beliefs ($\alpha^T = \frac{1}{2}$) will never strictly prefer to treat the workers asymmetrically rather than symmetrically (see condition (3)). Moreover, in an early period $t < T$, if the noise terms in all future periods are uniformly distributed, then Corollary 2 implies that if $\alpha^t = \frac{1}{2}$, $c^t = 0$ is optimal. We now show that the quality of the promotion decision will in some environments be increased by introducing bias in period $t < T$, even if $\alpha^t = \frac{1}{2}$. This possibility arises because in general (unlike the uniform case), optimal levels of bias in the future are sensitive to the current level. An interesting implication is that the organization may benefit from treating workers asymmetrically at the start of their careers, before any productive differences have been revealed. While the choice of individual to favour in the initial period could be made randomly, there would be no loss in efficiency from assigning the bias according to economically irrelevant demographic characteristics.

We analyse below the choice of first-period bias, but the results carry over to any period t when $\alpha^t = \frac{1}{2}$.

If $T = 1$, it is immediate from condition (3) that $c^1 = 0$ is optimal. In terms of the objective function $Q^1(c^1)$, which is the ex ante probability that the worker who wins the single contest is the more able, $Q^{1\prime}(0) = 0$. If the likelihood ratio $g^1(-c^1 - N)/g^1(-c^1 + N)$ is strictly decreasing at $c^1 = 0$, then $c^1 = 0$ is the unique optimum and $Q^{1\prime\prime}(0) < 0$. If the likelihood ratio is constant at 1 for $c^1 \in [-b^1, b^1]$, then any c^1 in this interval is optimal and $Q^{1\prime\prime}(0) = 0$.

For any $T \geqq 2$, the symmetric positions of the workers make the objective function symmetric about $c^1 = 0$, so $c^1 = 0$ must be a stationary point. Whether $c^1 = 0$ is a local minimizer or a local maximizer depends on the second-order effects of changes in c^1.

Consider the probability, α, assessed after the first period, that the winner of the first contest is better.[21] The larger is α, the more informative (ex post) was the first-round result. If $c^1 = 0$, the realized value of α is the same whoever wins the first round, whereas if $c^1 > 0$ (so i is favoured), the MLRC implies that α is larger if j wins than if i wins. Before the first contest, α is a random variable, whose mean and variance depend on c^1. Increasing $|c^1|$ from 0 (weakly) reduces the mean of α, since the mean of α is simply $Q^1(c^1)$, and $Q^{1\prime\prime}(0) \leqq 0$. But raising $|c^1|$ also increases the variability of α, and this change, ceteris paribus, is beneficial: the probability (labeled $V(\alpha)$) that the better worker is ultimately promoted, assessed at the beginning of period 2 and given that the biases in periods $2, \ldots, T$ are chosen optimally as functions of α, is a convex function of α. $V(\alpha)$ is strictly convex somewhere as long as the optimal levels of $c^2, \ldots, c^T$, given previous rank-order outcomes, are not independent of α. ($V(\alpha)$ is linear if $\Delta\varepsilon^t$ is uniformly distributed for $t \geqq 2$.) There are thus two opposing second-order effects of introducing bias in the first period. The negative effect is to reduce the mean level of informativeness of the first contest; this would be the only effect if the first period were also the last. The

21. The time superscript 2 is omitted for convenience.

positive effect is to make the informativeness of the first contest variable, since the decision-maker's flexibility to adjust future biases according to first-period informativeness induces a preference for such variability.

To compare these positive and negative effects, define the function $\alpha(z)$ as the probability that the winner of the first round is better, given that the bias "advantaged" him by z. (With first-period bias c^1 in favour of i, $\alpha = \alpha(c^1)$ if i wins and $\alpha = \alpha(-c^1)$ if j wins.) The function $\alpha(z)$ is decreasing, by (2), and is strictly decreasing at $z = 0$. Also define $A_i(c^1)$ as the probability that i wins the first round, given that the bias advantages i by c^1. The probability that, after period T, the decision-maker promotes the more able worker is

$$Q^T(c^1) = A_i(c^1)V(\alpha(c^1)) + (1 - A_i(c^1))V(\alpha(-c^1)).$$

Differentiating twice (assuming sufficient differentiability) and evaluating at $c^1 = 0$ gives

$$Q^{T''}(0) = V'(\alpha(0))(4A_i'(0)\alpha'(0) + \alpha''(0)) + V''(\alpha(0))(\alpha'(0))^2. \tag{4}$$

We can express $Q^1(c^1)$ as

$$Q^1(c^1) = A_i(c^1)\alpha(c^1) + (1 - A_i(c^1))\alpha(-c^1), \quad \text{so } Q^{1''}(0) = 4A_i'(0)\alpha'(0) + \alpha''(0). \tag{5}$$

Substitution of (5) into (4) shows that

Proposition 5. *$c^1 = 0$ is a local minimizer of the objective function for a T-period problem if and only if*

$$\frac{V''(\alpha(0))}{V'(\alpha(0))}(\alpha'(0))^2 > -Q^{1''}(0). \tag{6}$$

The right-hand side of (6) represents the (second-order) cost of the reduction in the mean level of first-period informativeness that accompanies an increase in $|c^1|$ from 0, and the left-hand side represents the (second-order) benefit from the increase in the variability of informativeness. Note the resemblance of the left-hand side to the Arrow–Pratt formula for a risk premium. Introducing bias in the initial contest is more likely to benefit the organization the larger the effect this would have on first-period informativeness, the more sensitive future optimal biases are to a change in informativeness, and the smaller the cost of bias in a single-period setting. The example in the next subsection is one in which biasing the first contest is optimal.

6.3. *Favouring a worker in an early period may reduce his promotion chances*

In the final period, the introduction of bias in favour of the leader always increases (at least weakly) the leader's chances of promotion, given an optimal promotion rule for the decision-maker. In an early period, however, the introduction of bias may reduce the promotion chances of the worker whom it favours in the short run, even when this bias improves the organization's decision. To understand this, note that, by the MLRC, introducing bias in the first period in i's favour lowers the decision-maker's degree of confidence in i whether he wins or loses the first contest. Therefore, if $T = 2$, Proposition 1 implies that the optimal second-period bias becomes less favourable to i, whatever the first-period outcome. If these adverse shifts in the second-period bias are sufficiently severe then, with respect to i's overall chance of promotion, they can outweigh the benefit from the increased likelihood of winning the first contest.

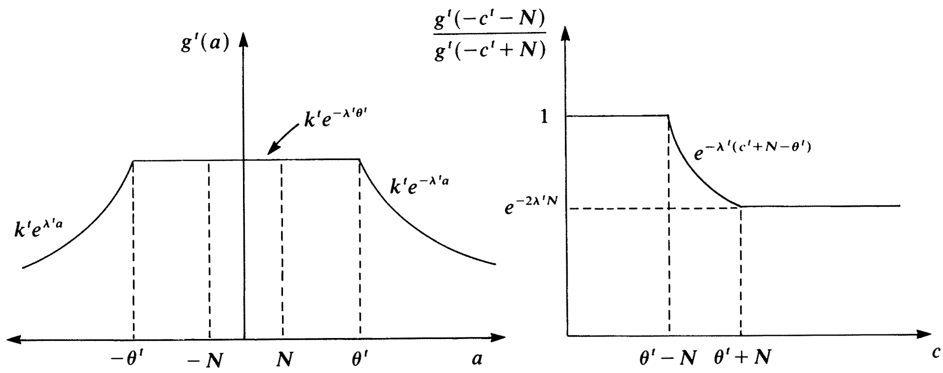

FIGURE 3

A two-period example in which introducing a small bias in favour of i in period 1 benefits the organization but reduces i's chances of promotion

Example 1. We present a two-period example in which the introduction of bias in i's favour in the first period benefits the organization but reduces i's promotion chances. Let the density functions $g^1(\cdot)$ and $g^2(\cdot)$ have exponential tails but flat centres (see Figure 3):

$$g^t(a) = \begin{cases} k^t \exp(-\lambda^t a) & \text{if } a > \theta^t \\ k^t \exp(-\lambda^t \theta^t) & \text{if } a \in [-\theta^t, \theta^t] \\ k^t \exp(\lambda^t a) & \text{if } a < -\theta^t. \end{cases}$$

Given (λ^t, θ^t), k^t is chosen to make $g^t(\cdot)$ a density. We confirm in the Appendix that there exists a set of parameters $(\lambda^1, \theta^1, \lambda^2, \theta^2)$ such that the following properties are satisfied: (1) a small increase in c^1 from 0 has no second-order effect (as well as no first-order effect) on the mean of α: $Q^{1\prime\prime}(0) = 0$; (2) $V(\alpha)$ is continuous, strictly convex for $\alpha < \alpha(0)$, and linear for $\alpha > \alpha(0)$; (3) when $c^1 = 0$, the optimal second-period bias in favour of the first-period winner is any value in $[\theta^2 + N, \infty)$, whereas when c^1 is increased slightly, the optimal bias in i's favour if i wins is slightly less than $\theta^2 + N$ and the optimal bias in j's favour if j wins is infinite. Properties (1) and (2) ensure that a small, though not infinitesimal, increase in c^1 from 0 benefits the organization, and property (3) implies that such a change reduces i's chance of promotion below $\frac{1}{2}$, the value that results from symmetric treatment of the workers when $c^1 = 0$.

7. THE USE OF BIAS WHEN THERE ARE MORE THAN TWO OPTIONS AFTER THE FINAL PERIOD

The analysis so far has assumed that, after the observation phase, the decision-maker faces only the binary choice of which worker to promote. What happens if he can also choose how much extra capital or responsibility to give the promoted worker? (In the spinner example, what if he can choose how much to bet, as well as on which colour?) Introducing a non-binary terminal decision raises the same issues for the choice of final-period bias as arise in the original model in choosing bias in early periods, since later decisions about levels of bias are not restricted to be binary.

When the decision-maker is restricted to choosing only which worker to promote, he sets the final-period bias to maximize his expected degree of confidence in the promoted worker, that is, his expected posterior probability that the promoted worker is better. On the other hand, in an early period, the flexibility to adjust future biases according to the informativeness of the current result means that he maximizes the expectation of a convex function of his degree of confidence in next period's leader. Similarly, when he is allowed to vary the treatment of the promoted worker according to his beliefs, he chooses the final-period bias to maximize the expectation of a convex function of his degree of confidence in this worker. In the latter two cases, the payoff is increasing in both the mean and the variability of the resulting degree of confidence.

Consequently, *when the decision-maker is not restricted to a binary decision after the observation phase, then even in the final period, i) he may strictly prefer to treat symmetrically rated workers asymmetrically and ii) when there is a leader, it may be optimal to set the bias against him.*

Example 2. Suppose that the decision-maker has three options after period T—promote i, promote j, or promote neither—and that, relative to making no promotion, the loss from promoting the less able worker exceeds the gain from promoting the more able one. Then there exists a critical value, $\bar{\alpha} \in (\frac{1}{2}, 1)$, of the degree of confidence after period T, α^{T+1}, above which the decision-maker will promote the worker with the better record and below which he will make no promotion. The objective function in period T is thus the expectation, with respect to $\Delta\varepsilon^T$, of $W(\alpha^{T+1})$, where W is constant for $\alpha^{T+1} < \bar{\alpha}$ and linearly increasing for $\alpha^{T+1} > \bar{\alpha}$.

First let $T=1$ and suppose that $\Delta\varepsilon^1$ is uniformly distributed. Assume that $\bar{\alpha}$ is sufficiently large that, if $|c^1| = E^1 - N$ and the favoured worker wins, the decision-maker strictly prefers to make no promotion. Then *the unique optimal magnitude of* c^1 *is* $|c^1| = E^1 - N$. By Proposition 3, $E(\alpha^2)$ (the objective function in the original model when $T=1$) is constant for $c^1 \in [-(E^1-N), E^1-N]$. However, by the MLRC, as $|c^1|$ is increased from 0 to $E^1 - N$, the value of α^2 if the favoured worker wins decreases, to a value less than $\bar{\alpha}$ (by assumption), while the value of α^2 if the disadvantaged worker wins increases, to 1. Given the shape of $W(\cdot)$, $EW(\alpha^2)$ is non-decreasing in $|c^1|$ and strictly increasing for $|c^1|$ less than but sufficiently close to $E^1 - N$. Increasing $|c^1|$ beyond $E^1 - N$ lowers $EW(\alpha^2)$, since the value of α^2 if the disadvantaged worker wins remains at 1, while the probability of this outcome falls.

Now let $T=2$ and suppose that $\Delta\varepsilon^1$ and $\Delta\varepsilon^2$ are uniformly distributed, with $E^1 = E^2$. Fix $c^1 = 0$ and suppose that i won the first round. Assume that $\bar{\alpha}$ is sufficiently large that, if $c^2 = 0$, then whatever the period-2 outcome, the decision-maker strictly prefers to make no promotion. Then *the optimal strategy is to set* c^2 *to disadvantage the leader, i, and, if i wins in period* 2, *to promote him, whereas if j wins, to promote neither.* To prove this, observe that, starting from $c^2 = 0$, the objective function can be increased by introducing a sufficiently large bias in favour of either i or j: a bias in i's (j's) favour can produce a sufficient increase in the degree of confidence in j (i) if j (i) wins to induce the decision-maker to promote j (i) (while promoting neither if i (j) wins). Favouring i by $c^2 = E^2 - N$ is strictly preferred to any other $c^2 \geqq 0$: as c^2 is increased from 0, $EW(\alpha^3)$ is constant until the decision-maker first becomes willing to promote j if j wins, and thereafter is increasing until $c^2 = E^2 - N$, since $E(\alpha^3)$ is increasing, $W(\cdot)$ is convex, and the value of α^3 if i (j) wins is decreasing (increasing). Exactly as for $T=1$, $EW(\alpha^3)$ falls as c^2 is raised above $E^2 - N$. However, favouring j by $E^2 - N$ is strictly preferred to favouring i by $E^2 - N$: in both cases, the decision-maker will make a promotion if and only if the disadvantaged worker wins, and will be certain that worker is better, but this outcome is more likely when the disadvantaged worker is the leader, i. Hence, the optimal strategy will use bias against the leader in the final round, but will promote him if he wins and no one otherwise.

APPENDIX

Proof of Proposition 2. Suppose i is the leader in period t, given the history R^{t-1}. The decision-maker's objective function is

$$
\begin{aligned}
&P(\text{promoting the better worker after } T \text{ periods} \mid R^{t-1}) \\
&\quad = P(\Delta\eta = N,\ i \text{ wins round } t \text{ with bias } c^t,\ i \text{ is promoted} \mid R^{t-1}) \\
&\qquad + P(\Delta\eta = -N,\ i \text{ wins round } t \text{ with bias } c^t,\ j \text{ is promoted} \mid R^{t-1}) \\
&\qquad + P(\Delta\eta = N,\ j \text{ wins round } t \text{ against bias } c^t,\ i \text{ is promoted} \mid R^{t-1}) \\
&\qquad + P(\Delta\eta = -N,\ j \text{ wins round } t \text{ against bias } c^t,\ j \text{ is promoted} \mid R^{t-1}) \\
&\quad = \alpha^t H^t(-c^t - N) P(i \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t,\ \Delta\eta = N) \\
&\qquad + (1-\alpha^t) H^t(-c^t + N) P(j \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t,\ \Delta\eta = -N) \\
&\qquad + \alpha^t H^t(c^t + N) P(i \text{ is promoted} \mid R^{t-1},\ j \text{ wins round } t,\ \Delta\eta = N) \\
&\qquad + (1-\alpha^t) H^t(c^t - N) P(j \text{ is promoted} \mid R^{t-1},\ j \text{ wins round } t,\ \Delta\eta = -N).
\end{aligned}
$$

Since the values of c^τ, $\tau = t+1, \ldots, T$, will be set optimally, given c^t, the first-order condition with respect to c^t is:

$$
\begin{aligned}
&\alpha^t g^t(-c^t - N) P(i \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t, \Delta\eta = N) \\
&\qquad + (1-\alpha^t) g^t(-c^t + N) P(j \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t, \Delta\eta = -N) \\
&\quad = \alpha^t g^t(c^t + N) P(i \text{ is promoted} \mid R^{t-1},\ j \text{ wins round } t, \Delta\eta = N) \\
&\qquad + (1-\alpha^t) g^t(c^t - N) P(j \text{ is promoted} \mid R^{t-1},\ j \text{ wins round } t, \Delta\eta = -N).
\end{aligned}
$$

Since

$$\frac{\alpha^t g^t(-c^t - N)}{(1-\alpha^t) g^t(-c^t + N)} = \frac{P(\Delta\eta = N \mid R^{t-1}, \Delta x^t = -c^t)}{P(\Delta\eta = -N \mid R^{t-1}, \Delta x^t = -c^t)},$$

the first-order condition can be written as

$$
\begin{aligned}
&P(i \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t, \Delta\eta = N) P(\Delta\eta = N \mid R^{t-1}, \Delta x^t = -c^t) \\
&\qquad + P(j \text{ is promoted} \mid R^{t-1},\ i \text{ wins round } t, \Delta\eta = -N) P(\Delta\eta = -N \mid R^{t-1}, \Delta x^t = -c^t) \\
&\quad = P(i \text{ is promoted} \mid R^{t-1}, j \text{ wins round } t, \Delta\eta = N) P(\Delta\eta = N \mid R^{t-1}, \Delta x^t = -c^t) \\
&\qquad + P(j \text{ is promoted} \mid R^{t-1}, j \text{ wins round } t, \Delta\eta = -N) P(\Delta\eta = -N \mid R^{t-1}, \Delta x^t = -c^t),
\end{aligned}
$$

or, equivalently,

$$
\begin{aligned}
&P(\text{better worker is promoted} \mid R^{t-1},\ i \text{ declared winner of round } t, \Delta x^t = -c^t) \\
&\quad = P(\text{better worker is promoted} \mid R^{t-1}, j \text{ declared winner of round } t, \Delta x^t = -c^t). \quad \|
\end{aligned}
$$

Proof of Proposition 4. *Sufficiency*: Let $\Delta\varepsilon^t$ be uniformly distributed on $[-E^t, E^t]$ for each $t \geqq 2$. If cardinal information on Δx^t is available in all periods, necessary and sufficient conditions for an optimal promotion policy are (B1), (B2), and (B3):

(B1) If for some $t \geqq 2$, $\Delta x^t \in I_1^t$, promote j.
(B2) If for some $t \geqq 2$, $\Delta x^t \in I_3^t$, promote i.

Let b^1 be the largest value of c^1 solving $g^1(-c^1 - N)/g^1(-c^1 + N) = 1$, the first-order condition for c^1 when $T = 1$ and $\alpha^1 = \frac{1}{2}$. By the symmetry of $g^1(\cdot)$, $-b^1$ is the smallest solution to this equation.

(B3) If for all $t \geqq 2$, $\Delta x^t \in I_2^t$, promote i if $\Delta x^1 > b^1$ and promote j if $\Delta x^1 < -b^1$.

A promotion policy satisfying (B1), (B2), and (B3) can be implemented, using only biased rank-order observations, by choosing $c^1 \in [-b^1, b^1]$ and for $t \geqq 2$ favouring the leader by $E^t - N$, where the leader is the winner of the previous round.

Necessity: Suppose first that the noise terms are uniformly distributed in all periods except τ, for some $\tau \geqslant 2$. The optimal policies for aggregating cardinal observations on Δx^t are independent of the order of the observations, so they must satisfy (B1), (B2), and (B3) (with the time superscripts appropriately modified). Therefore, if the decision-maker could implement one of these policies using bias, he would have to set $c^\tau \in [-b^\tau, b^\tau]$ (where b^τ is defined analogously to b^1 above). No matter how he sets the biases in other periods, there are realizations of $(\Delta x^1, \ldots, \Delta x^T)$ such that the biased rank-order observations in periods $t \neq \tau$ suggest, but do not conclusively establish, that the worker who lost the period-τ contest is more able. (When $\tau = 1$, this possibility never arises when biases are chosen optimally: if the observations in periods $t \geqq 2$ are not conclusive, they suggest that the period-1 winner is better.) The decision-maker with access only to biased rank-order observations cannot make the same promotion choice, for each such realization of $(\Delta x^1, \ldots, \Delta x^T)$, as when cardinal information on Δx^t is available; the optimal choice given cardinal information varies according to whether the realizations of Δx^t for $t \neq \tau$ are conclusive or uninformative.

Now suppose that there are at least two periods in which the noise terms have non-uniform distributions. We explicitly analyse the case $T = 2$. The argument extends in a straightforward manner to any $T > 2$. An optimal policy given cardinal information on Δx^1 and Δx^2 specifies, for each Δx^1, a cutoff value of Δx^2, above (below) which i (j) should be promoted. Given that $\Delta\varepsilon^1$ and $\Delta\varepsilon^2$ are not uniformly distributed and that the likelihood ratios in periods 1 and 2 are continuous for $\Delta x^1 \in (-(E^1 - N), E^1 - N)$ and $\Delta x^2 \in (-(E^2 - N), E^2 - N)$, respectively, there must exist an interval of Δx^1 values on which the cutoff value for Δx^2 is a continuous, strictly-decreasing function of Δx^1. However, when only a biased rank-order observation is available in the first period, a small change in Δx^1 either has no effect on the rank-order classification of the first-round outcome, and so has no effect on the second-period bias, or it changes the identity of the first-round winner, and so causes a discontinuous change in the second-period bias. Therefore, for Δx^1 values in this interval, the promotion choice when only biased rank-order observations are used cannot depend on Δx^2 in the same way as if cardinal information on Δx^1 and Δx^2 were available. ||

Proof of Corollary 2. First we establish that the sufficiency part of Proposition 4 holds for $\alpha^1 > \frac{1}{2}$ (as well as for $\alpha^1 = \frac{1}{2}$). Let i be the period-1 leader. In the proof of Proposition 4, (B3) must be replaced by

(B3′) If for all $t \geqq 2$, $\Delta x^t \in I^t_2$, promote i if $\Delta x^1 > -\underline{d}^1$ and promote j if $\Delta x^1 < -\bar{d}^1$, where $[\underline{d}^1, \bar{d}^1]$ is the (possibly degenerate) interval of values of c^1 solving $g^1(-c^1 - N)/g^1(-c^1 + N) = (1 - \alpha^1)/\alpha^1$, the first-order condition for c^1 when $T = 1$ and $\alpha^1 > \frac{1}{2}$. A promotion policy satisfying (B1), (B2), and (B3′) can be implemented, using only biased rank-order observations, by choosing $c^1 \in [\underline{d}^1, \bar{d}^1]$ and for $t \geqq 2$ favouring the leader by $E^t - N$, where the leader is the winner of the previous round. Thus the sufficiency part of Proposition 4 generalizes to $\alpha^1 \geqq \frac{1}{2}$. The strategy for setting bias just described is therefore optimal, for any T.

Now relabel period 1 as period t, period 2 as period $t+1$, etc. The previous argument then yields an optimal strategy from period t onwards, for any $\alpha^t \geqq \frac{1}{2}$ and any T, and therefore by condition (3), it is optimal to choose c^t as if t were the final period. ||

Proof of Corollary 3. From the proof of Corollary 2, if $\Delta\varepsilon^\tau$ is uniformly distributed for all $\tau > t$, then for any $\alpha^t \geqq \frac{1}{2}$, the decision-maker can achieve as high a probability of promoting the better worker, by using bias optimally in periods $t, t+1, \ldots, T$, as if he had access to cardinal information on Δx^t but only biased rank-order observations on Δx^τ for $\tau > t$. Therefore, a refinement in the information partition on Δx^t cannot increase the ex ante probability of promoting the better worker. ||

Proof of Corollary 4. Let $\bar{c}^1$ be an optimal value of c^1 when biased rank-order contests are used in both periods. Now let the decision-maker partition either of the two intervals of Δx^1 values, $(-\infty, -\bar{c}^1)$ or $(-\bar{c}^1, \infty)$, into two subintervals. We will show that for some choice of subintervals, the optimal value of c^2 must vary according to which subinterval Δx^1 was in. Therefore this refinement of the first-period partition is of strictly positive expected value. Since the optimal partition into three intervals must be at least as good as this refinement, the conclusion follows.

Case (a): Suppose $\bar{c}^1 > 0$ is an optimal first-period bias. (If a non-zero value of $\bar{c}^1$ is optimal, then by symmetry there is an optimal value that is positive.) Suppose the decision-maker partitions the interval $(-\bar{c}^1, \infty)$ into $S_1 \equiv (-\bar{c}^1, \bar{c}^1)$ and $S_2 \equiv (\bar{c}^1, \infty)$. If $\Delta x^1 \in S_1$, then his period-2 beliefs are $\alpha^2 = \frac{1}{2}$. If $\Delta x^1 \in S_2$, then he assigns probability $\alpha^2 > \frac{1}{2}$ to i being better. Since $\Delta\varepsilon^2$ is not uniformly distributed, the set of optimal c^2 values when $\alpha^2 = \frac{1}{2}$ is a (possibly degenerate) interval $[-b^2, b^2]$, where $0 \leqq b^2 < E^2 - N$. Since the likelihood ratio in period 2 is continuous and decreasing for $c^2 \in (-(E^2 - N), E^2 - N)$, any value of c^2 that is optimal for $\alpha^2 > \frac{1}{2}$ must be strictly greater than b^2.

Case (b): Suppose $\bar{c}^1 = 0$ is the unique optimal first-period bias. Suppose the decision-maker partitions the interval $(0, \infty)$ into $U_1 \equiv (0, y)$ and $U_2 \equiv (y, \infty)$. By the MLRC, he will assign a strictly higher probability to i being better if $\Delta x^1 \in U_2$ than if $\Delta x^1 \in U_1$: $\alpha^2_{(\Delta x^1 \in U_2)} > \alpha^2_{(\Delta x^1 \in U_1)}$. Since $\Delta\varepsilon^2$ is not uniformly distributed, the decision-maker can choose y sufficiently close to 0, and therefore $\alpha^2_{(\Delta x^1 \in U_1)}$ sufficiently close to $\frac{1}{2}$, so that the largest optimal value of c^2 for $\alpha^2_{(\Delta x^1 \in U_1)}$ (call this value b_y) is strictly less than $E^2 - N$. Since the likelihood ratio in period 2 is continuous and decreasing for $c^2 \in (-(E^2 - N), E^2 - N)$, any value of c^2 that is optimal for $\alpha^2_{(\Delta x^1 \in U_2)}$ must be strictly greater than b_y. ‖

Proof of Lemma 1. If j is declared the winner of the second round, then period-3 beliefs are $\alpha^3 = \frac{1}{2}$, and by condition (3), the optimal value of c^3 is 0. If i is declared the winner of the second round, then $\alpha^3 > \frac{1}{2}$. Since the likelihood ratio for $\Delta x^3 = -c^3$ never falls below $\exp(-2\lambda^3 N)$, condition (3) implies that the optimal value of c^3 is infinite if $((1-\alpha^3)/\alpha^3) < \exp(-2\lambda^3 N)$. Explicit calculation of α^3 for this case shows that as long as $\lambda^3 < (1/N)\ln(2\exp(\lambda^1 N) - 1) \equiv \bar{\lambda}^3(\lambda^1)$, where $\bar{\lambda}^3(\lambda^1) > \lambda^1$, the above inequality holds. ‖

Proof of Lemma 2. Section 5.1 shows that the sign of the first derivative of the objective function, at $c^2 = 0$, can be determined by identifying which worker the decision-maker, if he could observe a value $\Delta x^2 = 0$, would wish to declare the winner of round 2, given how the period-3 bias would depend on his period-2 report. If he would strictly prefer to declare j (i) the winner, then a first-order improvement can be achieved by a small reduction (increase) in c^2, favouring j (i).

Suppose that he observed $\Delta x^2 = 0$. If j were declared the winner of round 2, then since $c^3_{WL} = 0$, the probability that the better worker was ultimately promoted would be the probability that the better worker won the (unbiased) third round. If i were declared the second-round winner, then since $c^3_{WW} = \infty$ (for $\lambda^3 < \bar{\lambda}^3(\lambda^1)$), i would be certain to be promoted; the probability of promoting the more able worker would thus be the probability that i was better, given $\Delta x^1 > 0$ and $\Delta x^2 = 0$, which equals the probability that the better worker won the (unbiased) first round. Hence, given c^3_{WL} and c^3_{WW}, the decision-maker's preference of which worker to declare the second-round winner, if $\Delta x^1 > 0$ and $\Delta x^2 = 0$, would depend on the relative noisiness of the first- and third-round contests. If $\lambda^3 > \lambda^1$ (if $\lambda^3 < \lambda^1$), an unbiased contest in the third (first) round is more informative. Therefore, if $\bar{\lambda}^3(\lambda^1) > \lambda^3 > \lambda^1$ (if $\lambda^3 < \lambda^1$), the decision-maker would strictly prefer to declare j (i) the second-round winner, when $\Delta x^1 > 0$ and $\Delta x^2 = 0$. ‖

Analysis of Example 1. We show the existence of a set of parameters $(\lambda^1, \theta^1, \lambda^2, \theta^2)$ such that properties (1), (2), and (3) in the text hold. If $\theta^1 > N$, the likelihood ratio in period 1 is constant at 1 for $c^1 \in [-(\theta^1 - N), \theta^1 - N]$, so $Q^{1\prime\prime}(0) = 0$ and property 1) is satisfied.

Now suppose that θ^1 is such that $H^1(N)/H^1(-N) \equiv (1-\alpha(0))/\alpha(0) = \exp(-2\lambda^2 N)$. From Figure 3, if $c^1 = 0$, the optimal second-period bias in favour of the first-period winner is any value in $[\theta^2 + N, \infty)$. As long as the same size bias is used whichever worker wins, the ex ante probability that i is promoted is $\frac{1}{2}$ (by symmetry). Let c^1 be increased to a small positive value, δ. Since $(1-\alpha(\delta))/\alpha(\delta) > (1-\alpha(0))/\alpha(0) > (1-\alpha(-\delta))/\alpha(-\delta)$, the optimal second-period bias in i's favour if i wins will be slightly less than $\theta^2 + N$, whereas the optimal bias in j's favour if j wins will be infinite (so j will be promoted with certainty). Property (3) is thus satisfied. When $c^1 = \delta$, the ex ante probability that i is promoted equals the probability that i wins both contests, which as $\delta \to 0$ approaches

$$\tfrac{1}{2}[H^1(-N)H^2(-2N-\theta^2) + H^1(N)H^2(-\theta^2)] < \tfrac{1}{2}.$$

Therefore for $\delta > 0$ sufficiently small, i's probability of promotion is lower when $c^1 = \delta$ than when $c^1 = 0$.

From Figure 3, for $\alpha \in (\frac{1}{2}, \alpha(0))$, the optimal $|c^2|$ is continuously increasing in α, so $V(\alpha)$ is continuous and strictly convex in this range. For $\alpha > \alpha(0)$, the optimal $|c^2|$ is infinite, so $V(\alpha)$ is linear in this range. $V(\alpha)$ is continuous at $\alpha(0)$ by the Theorem of the Maximum. Property (2) is thus satisfied. Properties (1) and (2) imply that a small, though not infinitesimal, increase in c^1 from 0 increases the decision-maker's objective function.

The above arguments are valid for any $\theta^2 > 0$. However, we must ensure that the conditions $\theta^1 > N$ and $H^1(N)/H^1(-N) = \exp(-2\lambda^2 N)$ are compatible. Expressing $H^1(N)/H^1(-N)$ in terms of λ^1, θ^1 and N, equating to $\exp(-2\lambda^2 N)$, and solving for θ^1 shows that $\theta^1 > N$ and $H^1(N)/H^1(-N) = \exp(-2\lambda^2 N)$ are compatible if $2\lambda^1 N + 1 > \exp(2\lambda^2 N)$. Thus for any (λ^1, λ^2) satisfying $2\lambda^1 N + 1 > \exp(2\lambda^2 N)$, there exist (θ^1, θ^2) such that properties (1), (2), and (3) hold. ‖

Acknowledgment. I am very grateful to Jeremy Bulow, David Canning, David Cox, James Dow, Paul Klemperer, David Kreps, Jim Mirrlees, Dilip Mookherjee, John Moore, Canice Prendergast, Ariel Rubinstein, Barry Weingast, participants in seminars at Oxford, L.S.E., Tel Aviv, CORE, Stanford, and Marseille-Luminy, and two anonymous referees for helpful comments and discussions. Financial support and hospitality from the Stanford Institute for Theoretical Economics is gratefully acknowledged.

REFERENCES

AGHION, P., BOLTON, P., HARRIS, C. and JULLIEN, B. (1990), "Optimal Learning by Experimentation" (Discussion Paper 51, Nuffield College, Oxford).

ALPERN, S. and SNOWER, D. J. (1990), "Price-Quantity Decisions as Learning Instruments" (STICERD Discussion Paper TE/90/210, London School of Economics).

ARMITAGE, P. (1985), "The Search for Optimality in Clinical Trials", *International Statistical Review*, **53**, 15-24.

ARROW, K. (1974) *The Limits of Organization* (New York: W. W. Norton and Co.).

BALVERS, R. J. and COSIMANO, T. F. (1990), "Actively Learning About Demand and the Dynamics of Price Adjustment", *Economic Journal*, **100**, 882-898.

BARON, J. and BIELBY, W. (1986), "The Proliferation of Job Titles in Organizations", *Administrative Science Quarterly*, **31**, 561-586.

BERRY, D. A. and FRISTEDT, B. (1985) *Bandit Problems: Sequential Allocation of Experiments* (London: Chapman and Hall).

CALVERT, R. L. (1985), "The Value of Biased Information: A Rational Choice Model of Political Advice", *Journal of Politics*, **47**, 530-555.

DOW, J. (1991), "Search Decisions with Limited Memory", *Review of Economic Studies*, **58**, 1-14.

FISHMAN, M. and HAGERTY, K. (1990) "The Optimal Amount of Discretion to Allow in Disclosure", *Quarterly Journal of Economics*, **105**, 427-444.

GEANAKOPLOS, J. and MILGROM, P. (1988), "A Theory of Hierarchies Based on Limited Managerial Attention" (Cowles Foundation Discussion Paper 775R, Yale University, October, forthcoming in *Journal of the Japanese and International Economies*).

KANTER, R. M. (1977) *Men and Women of the Corporation* (New York: Basic Books).

KULKARNI, R. (1982), "Sequential Procedures, Adaptive", in Kotz, S. and Johnson, N. L. (eds.), *Encyclopedia of Statistical Sciences* (New York: John Wiley and Sons).

MEYER, M. (1986), "Incentives, Ability, and Career Profiles in Organizations" (Chapter 6, Ph.D. Thesis, Stanford Graduate School of Business).

RADNER, R. and STIGLITZ, J. (1984), "A Nonconcavity in the Value of Information", in Boyer, M., and Kihlstrom, R. (eds.), *Bayesian Models in Economic Theory* (Amsterdam: Elsevier).

RAIFFA, H. and SCHLAIFER, R. (1961) *Applied Statistical Decision Theory* (Boston: Graduate School of Business Administration, Harvard University).

ROSEN, S. (1986), "Prizes and Incentives in Elimination Tournaments", *American Economic Review*, **76**, 701-715.

ROSENBAUM, J. (1984) *Career Mobility in a Corporate Hierarchy* (New York: Academic Press).

ROTHSCHILD, M. (1974), "A Two-Armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, **9**, 185-202.

SAH, R. and STIGLITZ, J. (1986), "The Architecture of Economic Systems: Hierarchies and Polyarchies", *American Economic Review*, **76**, 716-727.

SIMON, H. (1976) *Administrative Behavior* (New York: Macmillan).

STIGLITZ, J. (1987), "Human Nature and Economic Organization" (Jacob Marschak Lecture, presented at the Far Eastern Meetings of the Econometric Society).

WICKLUND, R. A. and BREHM, J. W. (1976) *Perspectives on Cognitive Dissonance.* (Hillsdale, N. J.: Lawrence Earlbaum Associates).

ZELEN, M. (1969), "Play the Winner Rule and the Controlled Clinical Trial", *Journal of the American Statistical Association*, **64**, 131-146.