Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Modifying Variability and Correlations in Winner-Take-All Contests

Anil Gaba, Ilia Tsetlin, Robert L. Winkler,

OPERATIONS RESEARCH
Vol. 52, No. 3, May–June 2004, pp. 384–395
ISSN 0030-364X | EISSN 1526-5463 | 04 | 5203 | 0384

DOI 10.1287/opre.1030.0098

Modifying Variability and Correlations in Winner-Take-All Contests

Anil Gaba, Ilia Tsetlin
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676 {anil.gaba@insead.edu, ilia.tsetlin@insead.edu}

Robert L. Winkler
Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120, rwinkler@mail.duke.edu

We consider contests with a fixed proportion of winners based on relative performance. Special attention is paid to winner-take-all contests, which we define as contests with relatively few winners receiving relatively large awards, but we consider the full range of values of the proportion of winners. If a contestant has the opportunity to modify the distribution of her performance, what strategy is advantageous? When the proportion of winners is less than one-half, a riskier performance distribution is preferred; when this proportion is greater than one-half, it is better to choose a less risky distribution. Using a multinormal model, we consider modifications in the variability of the distribution and in correlations with the performance of other contestants. Increasing variability and decreasing correlations lead to improved chances of winning when the proportion of winners is less than one-half, and the opposite directions should be taken for proportions greater than one-half. Thus, it is better to take chances and to attempt to distance oneself from the other contestants (i.e., to break away from the herd) when there are few winners; a more conservative, herding strategy makes sense when there are many winners. Our analytical and numerical results indicate that the probability of winning can change substantially as variability and/or correlations are modified. Furthermore, in a game-theoretic setting in which all contestants can make modifications, choosing a riskier (less risky) performance distribution when the proportion of winners is low (high) is the dominant best-response strategy. We briefly consider some practical issues related to the recommended strategies and some possible extensions.

Subject classifications: decision analysis: strategies in contests, modifying variability, modifying correlations.
Area of review: Decision Analysis.
History: Received July 2001; revisions received April 2002, October 2002, February 2003; accepted March 2003.

1. Introduction

An increasing number of situations can be viewed as contests with a small number of winners based on relative performance rather than absolute performance, and with disproportionately large rewards going to the winners. We tend to think about contests most naturally in the realm of sports, where a golf tournament has only one winner and where a few superstars in some sports earn incomes (from endorsements and other activities as well as from their basic salaries) that dwarf the earnings of most of their fellow competitors. But contests occur in many other areas, whether or not they are actually billed as such. Consider, for example, the entertainment industry, where a small number of stars (actors, musicians, authors, television news anchors, etc.) command huge amounts relative to others. In academia, rankings of schools are watched closely; those ranking near the top are generally able to attract the best students and faculty and raise the most money. The faculty themselves engage in contests for appointments, for tenure, for endowed chairs, and at the most elite level, for prizes such as the Nobel Prize. In the business world, firms compete in research and development (R&D) races to develop new products, and employees compete for advancement, with huge rewards going to the firm that discovers the next wonder drug or the employee who rises to the position of CEO. Mutual funds compete for recognition in lists of top funds, which can lead to large inflows of money from investors and large rewards for the fund managers. Frank and Cook (1995) refer to contests with relatively small numbers of winners receiving relatively large rewards as *winner-take-all markets*, and claim (p. 7) "that the phenomenon has spread so widely and that so many of the top prizes have become so spectacular."

What sort of strategy is advantageous in a winner-take-all contest, and how might that strategy differ as the proportion of winners in the contest changes? A contestant faces a problem of decision making under uncertainty about her own performance as well as about the performance of others. Different strategies influence the distribution of her performance. All other things being equal, a simple location shift of the distribution to the right (where higher performance is better) is beneficial. Such a shift might be accomplished by means of increased effort; some work on the design of incentive schemes as contests has focused on effort level as a key decision variable (e.g., Lazear and Rosen 1981, O'Keeffe et al. 1984, Rosen 1986, Moldovanu and Sela 2001).

What about other aspects of the distribution, such as variability (dispersion)? March (1991) contrasts the exploration of new possibilities and the exploitation of old certainties by organizations and makes the following claim: "In competition to achieve relatively high positions, variability has a positive effect. In competition to avoid relatively low positions, variability has a negative effect" (p. 83). March gives an example where one contestant in a one-winner contest can change her mean and variance; the probability of winning is improved with higher means and/or higher variances. Inducement for risk taking (in terms of the spread of the output distribution) in Nash equilibria of one-winner contests is discussed in Dekel and Scotchmer (1999), Hvide (2002), Cabral (2003), and Hvide and Kristiansen (2003). On the empirical side, Sirri and Tufano (1998, p. 1590) find that "mutual fund consumers chase returns, flocking to funds with the highest recent returns, though failing to flee from poor performers." They suggest that funds could exploit this situation by increasing their variance of returns and hoping for an extraordinary return. Brown et al. (1996) and Chevalier and Ellison (1997) study incentives for mutual fund managers to increase or decrease the riskiness (variability) of their funds.

In a behavioral experiment, Gaba and Kalra (1999) find a strong tendency to increase variability when the proportion of winners in a sales contest is low. More generally, Gaba and Kalra consider characterizations of comparative risk in terms of mean preserving spreads (Rothschild and Stiglitz 1970, Pratt and Machina 1997). They study a contest where the performance levels of different contestants are independent and uniformly distributed and show that if the proportion of winners is low (high), a contestant prefers to have her performance distribution altered by a mean-preserving spread (mean-preserving contraction), thus attaining a riskier (less risky) distribution.

The work summarized above indicates that modifying variability of performance looks like a promising strategy for increasing the probability of winning in contests. This work, however, has assumed independence of contestants' performance levels and has not considered the possibility of modifying correlations with other contestants as a viable alternative to modifying variability. Also, general conclusions have been based primarily on distributions of performance that are uniform or two-outcome distributions, often for one-winner contests, and have focused on the direction of movement of the probability of winning, not on the magnitude of the change in probability.

In this paper, we consider the impact of modifying the distribution of performance on the probability of winning a contest with a fixed proportion of winners. For most of this paper, we take a decision-theoretic approach in which only one contestant is contemplating modifications, but we show that the resulting strategy is still optimal if all contestants can make modifications. In §2, we first generalize previous work by showing that increasing (decreasing) variability is optimal for all symmetric performance distributions when the proportion of winners is less (greater) than one-half and performance levels are independent. Then, we develop a model in which the performance levels of different contestants may be correlated; show that a riskier (less risky) performance distribution is preferred when the proportion of winners is less (greater) than one-half, where the distribution is affected by both variability and correlations; and derive an asymptotic formula for the probability of winning. To get an idea of the magnitude of improvements in the probability of winning due to modifying variability and/or correlation, we generate numerical results in §3 using a multinormal model. Our results indicate that increasing (decreasing) variability and reducing (increasing) correlations can considerably improve the probability of winning in contests with a low (high) proportion of winners. Increasing variability and reducing correlations can be viewed as two ways of attempting to break away from the herd, and we compare these two strategies in §3.3. In §4, we show that in a game-theoretic framework in which all contestants can modify their performance distributions, choosing a riskier (less risky) performance distribution is the dominant best-response strategy. We give a brief summary, along with a discussion of some practical issues and possible extensions, in §5.

2. Modifying Performance Distributions

We start with a brief summary of how §2 will unfold. First, we work with independent, symmetric performance distributions and show that a riskier (less risky) strategy is preferred when the proportion of winners p is less than (greater than) one-half. Then, we highlight the symmetry of concerns of contestants by showing that the probability of winning in a contest with proportion of winners p is equal to the probability of losing in a contest with proportion of winners $1-p$. Next, we develop a model with common and personal components of performance that allows correlated performance levels. Once again, a riskier (less risky) strategy is preferred when $p<(>)0.5$, and now the strategy can be affected by both variability and correlations. Finally, we derive a limiting value for the probability of winning that demonstrates a relationship between contests and quota-based systems.

Assume that a contest has $n \geqslant 2$ contestants $C_1, \ldots, C_n$, only one of whom (C_n) can modify her distribution of performance. The performance levels of all n contestants are independent with cumulative distribution functions (cdfs) F_i and probability density functions (pdfs) f_i, $i=1,\ldots,n-1$. C_n is choosing between cdfs F_n and G_n with corresponding pdfs f_n and g_n. The contestants with the k highest performance levels, $1 \leqslant k \leqslant n-1$, are winners in the contest, each receiving payoff W. Each of the remaining contestants receives payoff $L<W$. We assume that C_n wants to maximize her probability of winning and that she incurs no cost in choosing distributions. The proportion of winners in the contest, k/n, is denoted by p.

Definition. If the pdfs f and g with cdfs F and G are symmetric about some value m and $G(x) \leqslant F(x)$ for all $x > m$, then G is weakly riskier than F. If G is weakly riskier than F and $G(x) < F(x)$ for some $x > m$, then moving from F to G (from G to F) increases (decreases) riskiness.

We first show in Proposition 1 that riskier (less risky) performance distributions are preferred when the proportion of winners is less (greater) than one-half. When $p < 0.5$, on the one hand, the main concern is to reach a high performance level, which implies that the right tail is what matters and it is desirable to spread out or "thicken" that tail. When $p > 0.5$, on the other hand, the left tail is of concern, and it is desirable to shorten or "flatten" that tail. C_n will win the contest if her performance exceeds an order statistic of the performance levels of the other contestants, and we use a lemma (given in the appendix) concerning the density of this order statistic for independent symmetric distributions in the proof of Proposition 1.

Proposition 1. *If the performance levels are independent, $f_1, \ldots, f_n$ and g_n are symmetric about some value m, and G_n is weakly riskier than F_n, then C_n weakly prefers G_n to F_n if $p < 0.5$, weakly prefers F_n to G_n if $p > 0.5$, and is indifferent between G_n and F_n if $p = 0.5$.*

Proof. If $m \neq 0$, an equivalent contest with symmetry about zero can be obtained by means of a location shift, so without loss of generality we let $m = 0$. Let $T(x)$ represent the probability that C_n wins given that her performance is x; to simplify the notation, we suppress the dependence of $T(x)$ on n and $f_1, \ldots, f_{n-1}$. Because C_n wins if she beats $n-k$ of the other $n-1$ contestants, $T(x)$ is equal to the cdf of the $(n-k)$th order statistic of the $n-1$ performance levels of the other contestants, with corresponding pdf $t(x)$. A priori, before her performance level is known, C_n's probability of winning with F_n is $P_F(\text{win}) = \int_{-\infty}^{+\infty} T(x) f_n(x)\,dx = T(x)F_n(x)|_{-\infty}^{+\infty} - \int_{-\infty}^{+\infty} t(x)F_n(x)\,dx = 1 - \int_{-\infty}^{+\infty} t(x)F_n(x)\,dx$. Similarly, $P_G(\text{win}) = 1 - \int_{-\infty}^{+\infty} t(x)G_n(x)\,dx$, so $P_G(\text{win}) - P_F(\text{win}) = \int_{-\infty}^{+\infty} t(x)[F_n(x) - G_n(x)]\,dx$. From the lemma in the appendix, for all $x > 0$, $t(x) \geqslant t(-x)$ if $p < 0.5$, $t(x) \leqslant t(-x)$ if $p > 0.5$, and $t(x) = t(-x)$ if $p = 0.5$. Also, $F_n - G_n$ is an odd function with $F_n(x) - G_n(x) \geqslant 0$ for all $x > 0$. Thus, for $p < 0.5$, $t(x)[F_n(x) - G_n(x)] + t(-x)[F_n(-x) - G_n(-x)] \geqslant 0$ for all x. Therefore, $P_G(\text{win}) \geqslant P_F(\text{win})$. Similarly, $P_G(\text{win}) \leqslant P_F(\text{win})$ if $p > 0.5$ and $P_G(\text{win}) = P_F(\text{win})$ if $p = 0.5$. □

Corollary. *C_n strictly prefers G_n to F_n if $p < 0.5$ and there exists $x > 0$ such that $t(x) > t(-x)$ and $G_n(x) < F_n(x)$.*

The following proposition highlights the symmetry of concerns of contestants in $p > 0.5$ contests and $p < 0.5$ contests, showing that the probability of winning in a contest with proportion of winners p is equal to the probability of losing in a contest with proportion of winners $1 - p$. Without loss of generality, all symmetric pdfs in the rest of this paper are taken to be symmetric about zero.

Proposition 2. *Assume that the conditions of Proposition 1 are satisfied with $m = 0$. If $P_k(\text{win})$ denotes the probability that C_n wins the contest with k winners, then $P_k(\text{win}) = 1 - P_{n-k}(\text{win})$.*

Proof. Denote by t_{n-k} the pdf of the $(n-k)$th order statistic of the $n-1$ performance levels of the other contestants. Because all f_i are symmetric about zero, $t_{n-k}(x) = t_k(-x)$. Furthermore, $F_n(x) = 1 - F_n(-x)$. Thus, $P_k(\text{win}) = 1 - \int_{-\infty}^{+\infty} t_{n-k}(x)F_n(x)\,dx = 1 - \int_{-\infty}^{+\infty} t_k(-x)[1 - F_n(-x)]\,dx = \int_{-\infty}^{+\infty} t_k(x)F_n(x)\,dx = 1 - P_{n-k}(\text{win})$. □

Proposition 2 shows that a strategy that is beneficial (in terms of increasing the probability of winning) when the proportion of winners is p turns out to be harmful when that proportion is $1-p$. For example, a riskier performance distribution increases the probability of winning in a contest with $p < 0.5$ but decreases the probability of winning if $p > 0.5$. Proposition 2 enables us to focus on the $p < 0.5$ case in the discussion, because analogous conclusions hold in the opposite direction when $p > 0.5$. From the standpoint of a contest designer, it also shows that setting p low provides very different incentives from setting p high.

In many situations, the performance of different contestants may be correlated. A good example is a forecasting contest, because empirical evidence indicates relatively high positive correlations among forecasters (e.g., Figlewski and Urich 1983, Clemen and Winkler 1986). Different forecasters may have access to common data sets, similar training, and even information about the forecasts of others. In other contexts, positive correlations may be caused by the conditions under which the contest is held or by a tendency of the contestants to "stay with the herd."

We now develop a model that allows the consideration of correlations among contestants' performance levels. Suppose that for $i = 1, \ldots, n-1$, the performance of contestant i is given by $x_c + x_i$, where x_c has pdf f_c symmetric about zero and $x_1, \ldots, x_{n-1}$ are independent with pdfs f_i symmetric about zero. Thus, an individual's performance consists of a common component x_c shared with other contestants and a personal component x_i unique to the individual. The performance of C_n is of the form $ax_c + x_n$, where x_n is independent of x_c and $x_1, \ldots, x_{n-1}$ with pdf f_n symmetric about zero. This model allows for correlated performance levels through the common component (including negative correlations when $a < 0$) and allows C_n to modify her variability and/or her correlations with the others. The key to working with the model is to note that if we subtract x_c from everyone's performance, we get an equivalent contest (in terms of who wins) with independent performance levels $x_1, \ldots, x_{n-1}$ and $(a-1)x_c + x_n$.

Proposition 3. *Let f_n^* denote the pdf of $(a-1)x_c + x_n$ with corresponding cdf F_n^*. In the contest with correlated performance levels described above, C_n prefers a riskier F_n^* if $p < 0.5$, a less risky F_n^* if $p > 0.5$, and is indifferent if $p = 0.5$.*

Proof. Note that f_n^* is symmetric about zero. C_n's performance in the equivalent contest with independent performance levels has cdf F_n^*, and the other contestants have cdfs F_i, $i = 1, \ldots, n-1$. The result then follows from Proposition 1. □

Proposition 3 extends the result of Proposition 1 to a situation that allows correlated performance levels through a shared performance term. It is important to note that F_n^* can be affected by both variability and correlations, which means that modifying variability and modifying correlations are both reasonable strategies for C_n. This will become clearer in the context of the multinormal model in §3, where we express the roles of variability and correlations in F_n^* analytically.

The above results give directional strategies for increasing $P(\text{win})$, the probability of winning for C_n. What about the magnitude of changes in $P(\text{win})$? First, Proposition 4 gives a limiting value for $P(\text{win})$ that shows how the contest relates to a quota-based system. Then, in §3, we develop a multinormal model and use Proposition 4 to generate some numerical results.

Proposition 4. *Let $f_i = f$ for all $i \neq n$, with corresponding cdf F. Consider a sequence of contests such that $k/n \to p$ as $n \to \infty$. Then, $P(\text{win}) \to 1 - F_n^*[F^{-1}(1-p)]$ as $n \to \infty$, where F_n^* is as defined in Proposition 3.*

Proof. The probability that C_n wins in the equivalent independent contest is $P(\text{win}) = 1 - \int_{-\infty}^{+\infty} t(x) F_n^*(x)\, dx$, where $t(x) = [(n-1)!/((n-k-1)!(k-1)!)] f(x) F^{n-k-1}(x)[1 - F(x)]^{k-1}$. As $n \to \infty$ with p fixed, $t(x)$ becomes arbitrarily close to the delta function $\delta(x-q)$, where $q = F^{-1}(1-p)$, and $P(\text{win})$ therefore converges to $1 - F_n^*(q)$. In other words, the contest becomes equivalent to a quota-based compensation system for C_n with quota $q = F^{-1}(1-p)$. □

3. Modifying Variability and Correlations: A Multinormal Model

To get an idea of the magnitude of changes in $P(\text{win})$, we consider the case where the joint distribution of the contestants' performance levels is multinormal.

Multinormal Model. Suppose that the joint distribution of contestants' performance levels is multinormal with mean vector zero and covariance matrix V with elements $V_{ii} = 1\ \forall i \neq n$, $V_{nn} = s^{*2}$, $V_{ij} = r\ \forall i, j = 1, \ldots, n-1, i \neq j$, and $V_{in} = r^* s^*\ \forall i \neq n$. That is, C_n has standard deviation s^* and correlation r^* with all other contestants; the other $n-1$ contestants have standard deviation one and pairwise correlation r with each other.

As in §2, we use the model with common and unique components to transform this contest to an equivalent contest with independent performance levels. From C_n's perspective, the key parameter of interest in this transformed contest is $s^{*2} + r - 2r^*s^*$, which is the variance of the distribution of C_n's performance. This parameter, which is a function of s^*, r^*, and r, clarifies the roles of variability and correlations in finding an optimal strategy and influences $P(\text{win})$ for C_n in the context of the multinormal model. After presenting two general analytical results, we give some numerical examples to gain insight into what happens when C_n modifies variability and/or correlations. Although variability and correlations are intertwined in their influence on $P(\text{win})$, as shown in Propositions 5 and 6, we give numerical results separately for modifying variability and modifying correlations in §3.1 and §3.2, respectively, before considering the possibility of modifying both variability and correlations in §3.3.

Proposition 5. *For the multinormal model, C_n should maximize (minimize) $s^{*2} + r - 2r^*s^*$ if $p < (>)\ 0.5$.*

Proof. In the model with common and personal components with zero means from §2, let f_c be normal with variance r, $0 \leqslant r \leqslant 1$, f_i normal with variance $1 - r\ \forall i \neq n$, and f_n normal with variance $s^{*2} - a^2 r$, where $a = r^*s^*/r$. This is equivalent to the multinormal model given above, and f_n^* is normal with variance $s^{*2} + r - 2r^*s^*$. C_n's strategy then follows from the maximization or minimization of $s^{*2} + r - 2r^*s^*$ implied by Proposition 3. □

Proposition 6. *Under the conditions of Proposition 5, the probability of winning for contestant C_n approaches $1 - \Phi[(1-r)^{1/2}\Phi^{-1}(1-p)/(s^{*2} + r - 2r^*s^*)^{1/2}]$ as $n \to \infty$, where Φ is the standard normal cdf.*

Proof. This result follows directly from Proposition 4. □

We use this large-n result to calculate $P(\text{win})$ in the figures throughout the paper. To investigate the sensitivity of these results to variations in n, we calculated $P(\text{win})$ for various finite values of n and reproduced all of the figures for $n = 10$. The changes in the probabilities from the limiting case to $n = 10$ are surprisingly small and the relative changes are even smaller, so the numerical results for the large-n case are representative of virtually all values of n.

3.1. Modifying Variability

To study the impact of modifying variability in the context of the multinormal model, we first show the strategy that should be followed by C_n and then present some numerical results to study the magnitude of the changes in $P(\text{win})$ as C_n modifies variability.

Proposition 7. *For the multinormal model with r^* fixed, C_n should choose s^* as large as possible if $p < 0.5$, equal to r^* if $p > 0.5$, and at any level if $p = 0.5$.*

Proof. From Proposition 5, C_n should maximize (minimize) $s^{*2} + r - 2r^*s^*$ if $p < (>)\ 0.5$, which means maximizing s^* if $p < 0.5$ and setting $s^* = r^*$ if $p > 0.5$. If $p = 0.5$, s^* does not affect $P(\text{win})$. □

Figure 1 shows $P(\text{win})$ as a function of the modified standard deviation s^* for $p = 0.1(0.2)0.9$ and $r = r^* = 0$, 0.5, and 0.9. As implied by Proposition 2, the curves for

Figure 1. Impact of modifying variability: $P(\text{win})$ as a function of s^* for $p = 0.1(0.2)0.9$ and $r = r^* = 0$, 0.5, and 0.9.

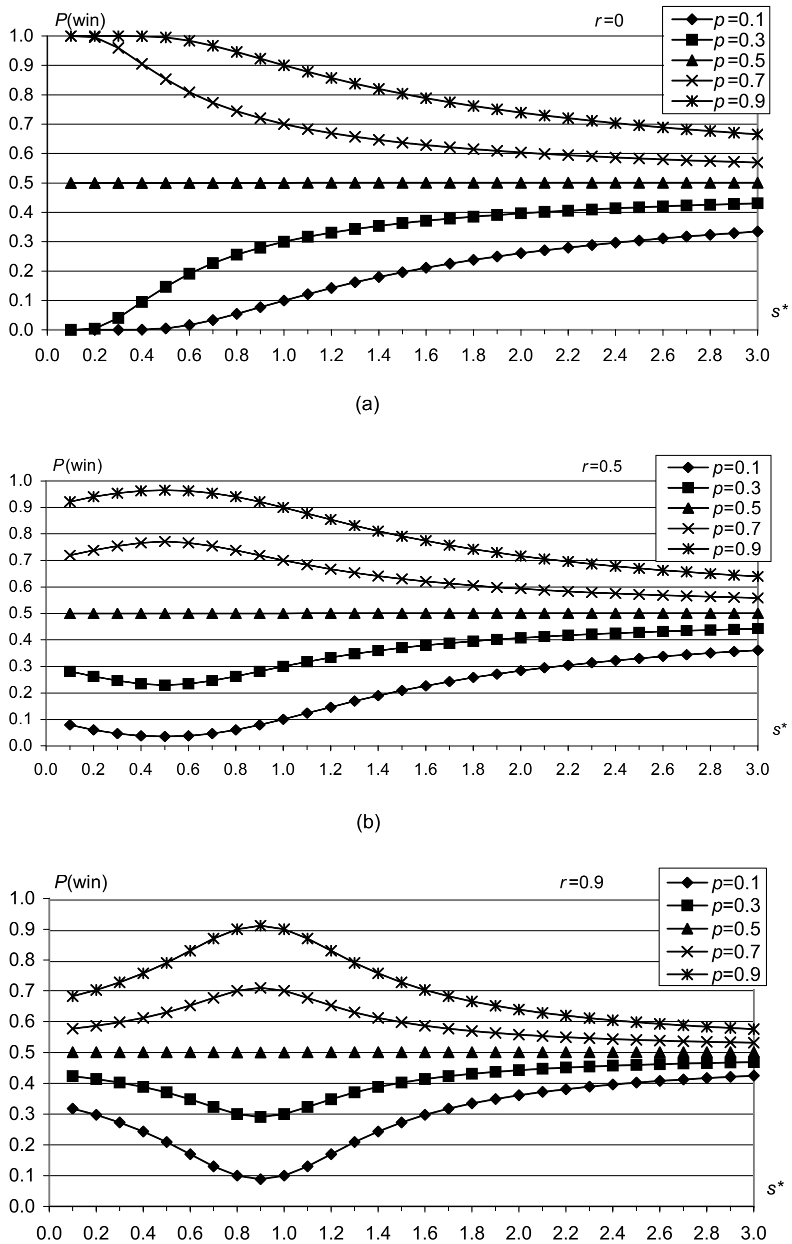

$p < 0.5$ are mirror images of the curves for $p > 0.5$. In the former case, $P(\text{win})$ is an increasing function of s^* when $r = 0$, which is as expected. The impact on $P(\text{win})$ of increasing s^* in contests with low proportions of winners and decreasing s^* in contests with high proportions of winners can be substantial. To simplify the discussion, we will focus on the $p < 0.5$ case. When $p = 0.1$ and $r = 0$, doubling s^* from the base rate of 1 increases $P(\text{win})$ by a factor of 2.7, whereas cutting s^* in half reduces $P(\text{win})$ by a factor of 20. When $p = 0.3$ and $r = 0$, the impact on $P(\text{win})$ of moving from the base case is smaller, increasing by a factor of 1.3 when doubling s^* and decreasing by a factor of 2 when halving s^*. Of course, as p gets closer to 0.5, there is less room for improvement because modifying variability cannot push $P(\text{win})$ above its limiting value of 0.5.

What about the impact of correlated performance levels? For $s^* > 1$, $P(\text{win})$ increases as s^* increases. Moreover, $P(\text{win})$ increases as the base correlation r increases for a given $s^* > 1$. That is, stronger positive correlations make increasing s^* even more attractive, particularly when r is quite high. For $s^* < 1$, however, the curves in Figures 1b and 1c exhibit nonmonotonic behavior, with the slopes changing sign at $s^* = r$. This nonmonotonic behavior might seem strange at first glance, but if the correlations are sufficiently high, the probability that different performance levels are on the same side of the mean increases. When they are above the mean (i.e., they are positive), the performance level from a contestant with higher standard deviation is more likely to be in the winning range. In the extreme, with perfect correlations of 1, the contestant with the highest standard deviation will have the highest performance level with probability 0.5 (i.e., when the performance levels are above zero). This is consistent with the results in Winkler and Brooks (1980) that in competitive bidding with two bidders who bid their best estimates, the "winner's curse" can be avoided by the bidder with the lower error variance if the correlation between their errors is higher than the ratio of the smaller to the larger error standard deviation. For $s^* < 1$, this ratio is $s^*/1 = s^*$, which is consistent with the change in the sign of the slopes at $s^* = r$ (both numerically from Figure 1 and analytically from Proposition 6 when $r^* = r$). From Figure 1, reducing s^* sufficiently can lead to substantial improvements in $P(\text{win})$ when the proportion of winners is low and r is high. However, C_n is still better off increasing variability when p is low, and in many contests it might be easier to increase variability than to reduce variability while holding the mean constant, especially when r is high.

These results are intuitively reasonable. As correlations get higher, the performance levels will be more tightly clustered, all else being equal, and a given increase in variability has a greater chance of moving C_n "outside the cluster." In turn, being outside the cluster is even more valuable as the proportion of winners in the contest is reduced. When s^* is so large that C_n is virtually assured of having an extreme value, $P(\text{win})$ is arbitrarily close to 0.5.

3.2. Modifying Correlations

In §3.2, we focused on the impact of modifications in variability by C_n. But because correlations can influence the probability of winning, as shown in Propositions 5 and 6, perhaps C_n could benefit by modifying correlations. For instance, when there are very few winners, it may be beneficial to break away from the herd by modifying one's correlations with other contestants. A fund manager might adopt a contrarian strategy, a sailor could take a completely different tack from that taken by other sailors in a boat race, a scientist trying to develop a new drug could try a radically different paradigm, and a forecaster could use a nonstandard methodology to develop forecasts.

Proposition 8. *For the multinormal model, C_n should choose r^* as small as possible if $p < 0.5$, as large as possible if $p > 0.5$, and at any level if $p = 0.5$.*

Proof. From Proposition 5, C_n should maximize (minimize) $s^{*^2} + r - 2r^*s^*$ if $p < (>)\ 0.5$, which means minimizing (maximizing) r^*. If $p = 0.5$, r^* does not affect $P(\text{win})$. □

To consider the magnitude of changes in $P(\text{win})$, we consider the case with all standard deviations (including C_n's) equal to 1 to isolate the influence of modifications in correlations. Figure 2 shows $P(\text{win})$ as a function of the modified correlation r^* for base correlations $r = 0.5$ and 0.9 and selected values of p. Note that the set of feasible values for r^* depends on r because of the constraint that the overall correlation matrix be positive definite. The $r = 0$ case is not shown because r^* can deviate from r only by $(n-1)^{-1/2}$ (which goes to zero as n gets large) when $r = 0$. Once again, there is symmetry between the results for p and $1-p$, and we will focus on $p < 0.5$. As implied by Proposition 6, $P(\text{win})$ decreases as r^* increases when $p < 0.5$. Looking at the same results from a different vantage point (graph not shown), we see that $P(\text{win})$ is an increasing function of r for given r^*.

The results support the intuition that reducing one's correlations with the other contestants can improve the chances of winning. The flexibility in modifying one's correlations and the gains from reducing the modified correlation r^* below r are limited when performance levels are independent, but greater when the contestants are highly correlated. For example, with $r = 0.9$ and $p = 0.1$, reducing r^* from 0.9 to 0.5 increases $P(\text{win})$ from 0.10 to 0.34 and reducing it further to zero (-0.5) moves $P(\text{win})$ to 0.38 (0.41). When $r = 0.5$ and $p = 0.1$, reducing r^* from 0.5 to zero (-0.5) increases $P(\text{win})$ from 0.10 only to 0.23 (0.28).

3.3. Modifying Variability vs. Modifying Correlations

We have seen that it is possible to increase the chance of winning a contest by modifying variability as in §3.1 (increasing variability when p is low, decreasing it when p is high) or correlations as in §3.2 (reducing correlations when p is low, increasing them when p is high). Given a choice between these two strategies, which seems more effective? Also, is it beneficial to combine strategies of modifying variability and correlations in some way? Note that doing something to modify correlations may be likely to modify variability in some manner, as well. For example, when a sailor takes a different tack from the rest of the sailors in a boat race, variability will probably change and correlations will probably decrease.

The impacts of changes in variability and correlations when $p = 0.1$ are shown in Figure 3 for three base cases: $r = 0.1, 0.5$, and 0.9. Changes in variability are represented by the horizontal axis, which is s^*, and changes in correlations are indicated by the different curves for different values of r^*. In each of the three graphs in Figure 3, the vertical axis measures $\Delta P(\text{win})$, the change in the probability of winning as compared with the base case (where $s^* = 1$ and $r^* = r$).

In a way, comparing changes in variability with changes in correlations is like comparing apples and oranges, because they are different measures. We would need to know the "costs" of changes to be able to say that one

Figure 2. Impact of modifying correlations: $P(\text{win})$ as a function of r^* for $p = 0.1(0.2)0.9$ and $r = 0.5$ and 0.9.

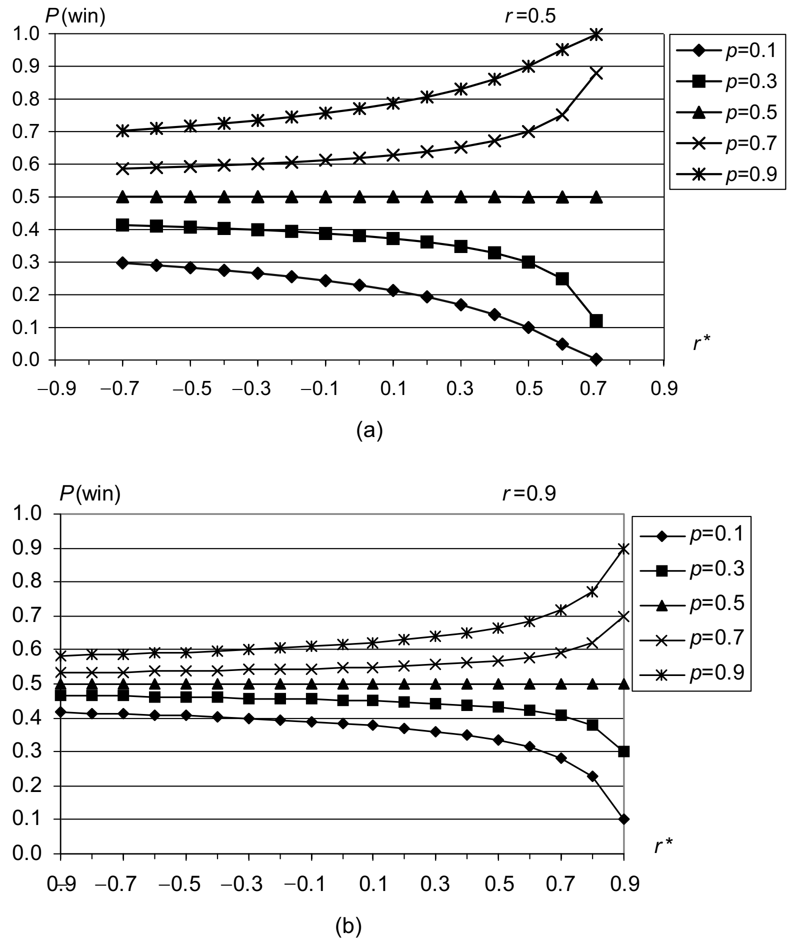

change is preferable to another. However, Figure 3 can give us an idea of what changes in these measures are needed for certain levels of improvement in the chance of winning. For example, when $r = 0.5$ an improvement of roughly 0.10 in $P(\text{win})$ can be attained by increasing the standard deviation to $s^* = 1.45$, by decreasing the correlations to $r^* = 0.2$, or by a combination of $s^* = 1.2$ and $r^* = 0.3$. As the base correlation r gets larger (i.e., as we move from Figure 3a to 3b and 3c), it takes a smaller increase in s^* or a smaller decrease in r^* to achieve a given $\Delta P(\text{win})$, and greater gains are possible for a given change in variability or correlations. For example, increasing the standard deviation to $s^* = 1.5$ improves $P(\text{win})$ by 0.10 when $r = 0.1$, by 0.11 when $r = 0.5$, and by 0.17 when $r = 0.9$; decreasing the correlations by 0.4 improves $P(\text{win})$ by 0.08 when $r = 0.1$, by 0.11 when $r = 0.5$, and by 0.24 when $r = 0.9$. Decreasing correlations appears to be relatively more effective for high r, whereas increasing variability is relatively more effective for low r.

If we look at modifying both variability and correlations, the gains are generally subadditive. On the one hand, when $r = 0.9$, the improvement in $P(\text{win})$ from increasing the standard deviation s^* to 1.5 is 0.17, from decreasing the correlation r^* to 0.5 is 0.24, and from doing both is $0.28 < 0.17 + 0.24$. On the other hand, once r^* is decreased to 0.5, an even larger decrease to $r^* = 0$ moves $\Delta P(\text{win})$ only to 0.28, exactly where it would be if r^* were left at 0.5 and s^* increased to 1.5. Similarly, if $r^* = 0.5$, moving s^* to 2 leads to about the same improvement (to $P(\text{win}) = 0.31$) as pushing r^* down even further to -0.5. The marginal gains tend to decrease as s^* increases or r^* decreases, leading to effects similar to the subadditivity when both s^* and r^* are changed.

In Figure 4, we fix $r = 0.5$ and $r^* = 0$ (i.e., C_n has already modified the correlations) and look at $\Delta P(\text{win})$ as a function of the modified standard deviation s^* for different values of p. For $p = 0.1$, the marginal gains when modifying s^* from 1 to 1.5 and 2 are 0.06 and 0.10, respectively. From Figure 1b, these gains are 0.09 and 0.16 when $r = 0.5$ and $r^* = 0.5$ (no change in correlations by C_n), supporting the notion that once modifications have increased the chance of winning, further modifications are less effective. Similar results hold for $p = 0.3$, and we see from Figure 4 that the gains from $p = 0.3$ are similar in nature but smaller in magnitude than those for $p = 0.1$.

Figure 3. Modifying variability and correlations: Change in P(win) as a function of s^* for $p = 0.1$, $r = 0.1(0.4)0.9$, and selected values of r^*.

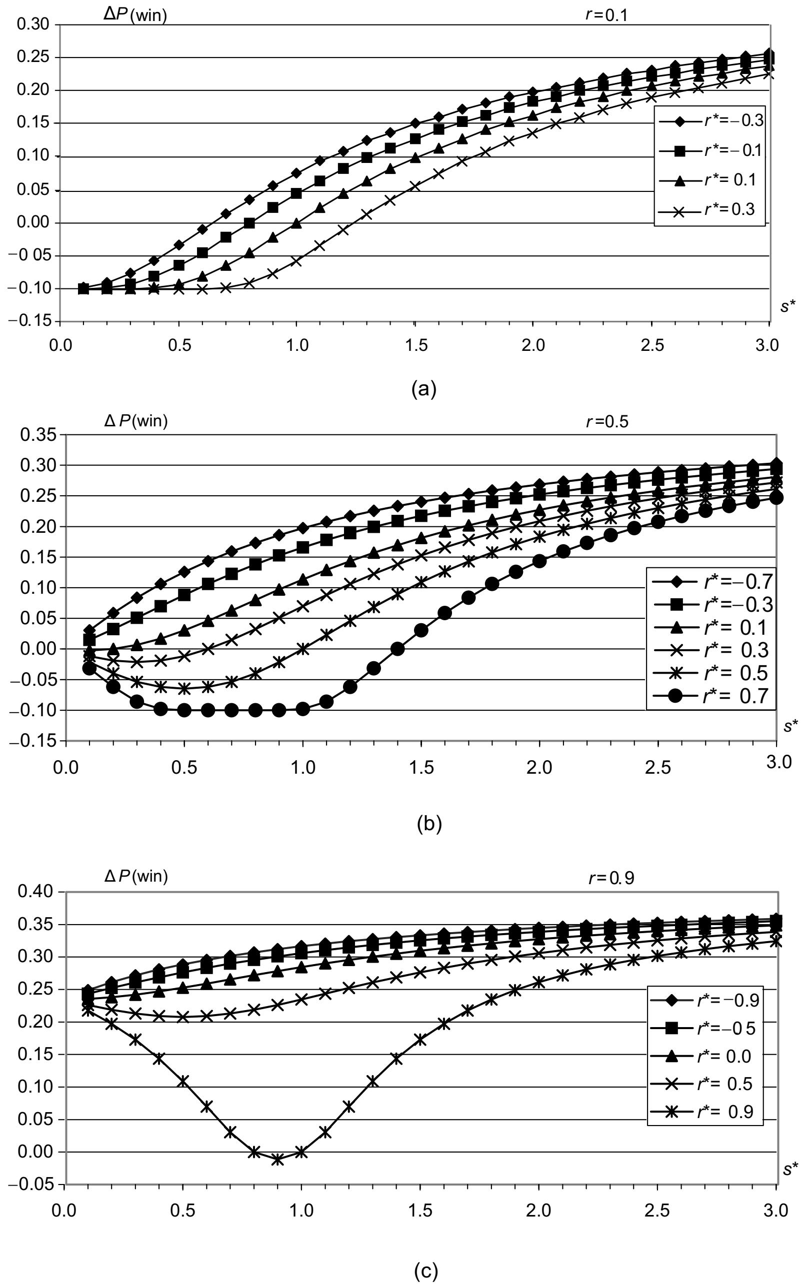

4. A Game-Theoretic Framework

We have considered a decision-theoretic framework from the viewpoint of C_n. As a result, we have not taken game-theoretic issues into consideration. In some cases, C_n might have a special opportunity not available to others. If not, other contestants might also be contemplating modifications. Here, we show that the results of §2 extend to the case where all contestants can modify their performance distributions.

Much of the game-theoretic work on contests in the economics literature has focused on a contestant's effort level as a decision variable; see Moldovanu and Sela (2001) for recent results and extensive references to related streams of work. Hvide (2002) considers both effort level and risk (variability) as decision variables in a contest with one winner. He analyzes the two-contestant case with independent performance levels having unimodal and symmetric distributions and shows that maximizing variance and minimizing effort is the Nash equilibrium. Hvide and Kristiansen (2003) develop a discrete model with two possible strategies and two types of competitors, focusing on selection efficiency. Cabral (2003) develops an infinite-period model

Figure 4. Change in P(win) as a function of s^* for $r = 0.5$, $r^* = 0$, and $p = 0.1(0.2)0.9$.

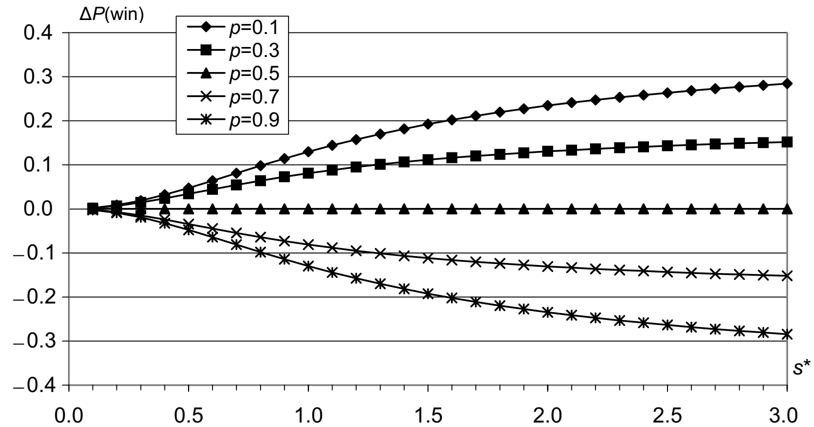

with two players and two alternatives (safe and risky) in each period. He provides conditions under which the competitor who is behind (ahead) should choose the risky (safe) strategy and conditions under which the competitor who is behind (ahead) prefers to take the opposite (same) strategy as that taken by the other competitor. Proposition 9 is consistent with these findings and extends them to contests with multiple winners, arbitrary (but symmetric) performance distributions, and correlated performance levels. However, we assume a single-period setting and do not include effort as a decision variable.

Proposition 9. *Suppose that contestant C_i, $i = 1, \ldots, n$, has performance level $x_c + x_i$, where $x_1, \ldots, x_n$, and x_c are independent. There are no restrictions on the distribution of x_c, and C_i has a choice among distributions for x_i from the set Ψ_i. All distributions in Ψ_i are symmetric about zero, and each Ψ_i contains cdfs G_i and F_i such that for any cdf F from Ψ_i, F is weakly riskier than F_i and G_i is weakly riskier than F. Then, the weakly dominant strategy for each contestant C_i is to choose G_i (the most risky cdf) if $p < 0.5$ and F_i (the least risky cdf) if $p > 0.5$.*

Proof. As in Proposition 3, the contest is equivalent to one with independent performance levels $x_1, \ldots, x_n$. Suppose that all contestants C_j, $j \neq i$, have chosen their distributions from Ψ_j. From Proposition 1, the best reply for C_i is to choose G_i if $p < 0.5$ and F_i if $p > 0.5$. □

In this situation, our recommended strategy from Proposition 3 is still best for C_n even though C_n's chance of winning would decrease as more competitors exercise their choice of strategies. From a decision-making viewpoint, then, the game-theoretic framework does not change much. Because Proposition 9 gives the dominant strategy, the strategy itself does not depend on whether the other contestants modify their distributions or on whether their distributions are known to the decision maker. Note that if the conditions of the corollary to Proposition 1 hold, the n-tuples $(G_1, \ldots, G_n)$ and $(F_1, \ldots, F_n)$ are the unique Nash equilibria for $p < (>)\ 0.5$.

Perhaps not all contestants are able to increase their variability or decrease their correlations, however, or perhaps not all who are able will choose to do so. There may be powerful psychological factors that discourage contestants from trying to break away from the herd. The strategy of increasing variability means that the probability of a very high performance is increased. It also means that the probability of a very low performance is increased, so there is a good chance that the contestant will be at the bottom of the heap. If the only outcomes that really matter are winning or not winning, then the possibility of being at the bottom should be irrelevant. However, in a psychological or social sense it might be important, especially when connected with a strategy that differs from what other contestants are doing and, by extension, from what is the commonly accepted strategy for the contest. Also, although increasing variability is advantageous in the context of a contest, it may be perceived by some contestants as undesirable. If some but not all other contestants modify variability or correlations, that would most likely dilute the increase in C_n's probability of winning compared with the case where no one else modifies, but still leave C_n better off with the modifications, as implied by the dominant strategy of Proposition 9.

5. Summary and Discussion

In contests based on relative performance, differentiating oneself from other contestants may improve one's chance of winning. Our results show that choosing a riskier (less risky) performance distribution is a good strategy when the proportion of winners p in the contest is less than (greater than) one-half. In the case of a multinormal joint performance distribution, this strategy can be implemented by modifying the variability of one's performance and/or the correlations of one's performance with the performance of other contestants. When the proportion of winners in the contest is less than one-half, increasing variability and decreasing correlations lead to improved chances of winning. In some ways these might be viewed as risky strategies, and indeed they do increase the chance of an especially poor performance. But if the only goal is to be one of the winners of the contest, and losing by a lot is no worse than losing by a little, the lower tail of the distribution of performance is irrelevant. McCardle and Winkler (1992) show that a strategy involving a riskier distribution might be preferred in the context of repeated gambles. Here, a riskier distribution is preferred in the context of a contest with few winners.

Of course, if $p > 0.5$, then the proper advice is the opposite: decrease variability and increase correlations. Our discussion focuses on the case with $p < 0.5$, but the symmetry of the situation (Proposition 2) means that similar implications hold for high p with the opposite directions of change being optimal.

In §2, we show that for symmetric performance distributions, the opportunity to change one's performance distribution is beneficial both in an independent world and in a correlated world. Because this distribution is influenced by both variability and correlations with other contestants, these results have implications for modifications in both of these aspects of the distribution. Our numerical results with the multinormal model in §3 then demonstrate that the probability of winning can change substantially as variability and/or correlations are modified, and show the magnitude of such changes as a function of the proportion of winners in the contest, the correlations among different contestants, and the degree to which variability and/or correlations are modified. Increasing variability leads to greater improvements in the chance of winning when correlations are high, with these improvements especially pronounced as the proportion of winners gets smaller. When correlations are high, breaking away from the herd in some manner can be even more beneficial, and increasing variability is one way of doing so.

Another way of breaking away from the herd is to reduce one's correlations with the other contestants. The results in §3.2 demonstrate that reducing correlations can indeed increase the probability of winning, with the greatest potential benefits coming when the initial correlations are quite high.

Comparisons of modifying variability and modifying correlations in §3.3 indicate that decreasing one's correlations with other contestants appears to be relatively more effective in improving $P(\text{win})$ for high initial correlations and increasing variability is relatively more effective for low initial correlations. Of course, without considering the feasibility and the cost (in money, time, effort, etc.) of these changes, it is hard to say which strategy might work best in a given situation. In general, the marginal improvements due to additional modifications decrease as the extent of the modifications already undertaken increases. In a similar fashion, increases in the chance of winning from a combination of modifying variability *and* modifying correlations are subadditive.

If C_n finds it advantageous to modify her performance distribution, then it is likely that other contestants will also arrive at the same conclusion. Therefore, it is reassuring to know that our results extend to the game-theoretic framework where all contestants have the opportunity to make modifications. In §4, we show that choosing a riskier (less risky) performance distribution when the proportion of winners is low (high) is the dominant best-response strategy.

Our model is simple but should be reasonably robust. Changes in the form of the distribution of performance would change the numbers but are not likely to change the general picture. When p is small, the right-hand tail of the distribution is what really matters, and when p is large, the left-hand tail drives C_n's strategy. Given the importance of the tails, a promising extension for further study involves the possibility of modifications to skewed distributions.

Some contests have a fixed number of winners k instead of a fixed proportion of winners. As the number of contestants n increases, $p = k/n$ decreases. We have seen that as p becomes smaller, the gains in $P(\text{win})$ associated with increased variability or decreased correlations tend to be larger. Thus, these strategies should be more effective as n increases in a contest with a fixed number of winners, all other assumptions remaining unchanged.

The assumptions of common means, variances, and correlations seem somewhat restrictive but might not be far off the mark for situations where the contestants are all experts in their fields, or where preliminary contests have already narrowed down the set of contestants. In many cases (e.g., sales contests), effort can influence results; our model implies that the contestants are all putting forth more or less the same (e.g., maximum) effort. In reality the means might differ somewhat, through more effort or better ability (working harder or smarter). However, with real experts or a selected group of contestants, the differences in means might often be quite small relative to the variability of the performance distributions. An interesting extension for future study would be to assume a hierarchical model with distributions of the model parameters among the contestants and a multiattribute utility function for C_n with effort as one attribute and winning or losing as another.

It is one thing to suggest that modifying one's performance distribution is a good strategy to follow in contests, and yet another thing to implement that strategy. Increasing variability may often be possible. For example, a salesperson can choose to pursue customers for whom the probability of getting an order is small but the size of the potential order is large. A mutual fund manager can increase the riskiness of a fund's portfolio. An academic in a tenure contest can work on problems that are acknowledged to be important but especially difficult. Similarly, reducing correlations by breaking away from the herd and taking a different approach is often feasible. The salesperson could try to sell to new customers who have never been approached before, the mutual fund manager can act in a "contrarian" manner, and the academic can try to take an entirely new approach to some problem rather than just adding another study similar to those that have been done.

The degree to which one can make such modifications and do so without incurring a reduction in mean performance or some other cost will depend on the nature of the contest. Of course, if the modifications are substantial enough, they are likely to overcome the effect of a reduction in the mean or other cost. March (1991, p. 85) calls a contest with low p "a right-hand-tail 'race' in which average performance (due to ability and effort) becomes irrelevant."

The structure of the contest implies that there are only two possible outcomes: winning and not winning. This means that there is no distinction among winners, nor is there any advantage to being near the top of the nonwinners in terms of performance. This is precisely the case in many sporting contests and is very close to reality in situations such as a research and development contest to develop a new drug. In other situations, such as a tenure contest for faculty, a partnership contest for consultants, or a sales contest, the structure does not fit perfectly but still is close enough to have implications for desirable strategies.

The strict dichotomy implies that utility for one's proportional rank among the contestants is a step function with a single step (down) at p. All contests with serious prizes for winning have a large step at p, not only because of the immediate rewards but also because of anticipated future benefits. In practice, contestants must consider complex situations with some issues peripheral to but not totally unrelated to the contest. Strategies may be affected by these issues, which could include some negative ramifications of lower performance within the losing range. This could cause the utility for proportional rank higher (worse) than p to be decreasing rather than constant, but the rate of decrease may be relatively low, in which case the step at p is still the dominant factor. Contests with another step or a steep slope at high proportional ranks (low performance levels) would not fit the structure in this paper well because the increased chance of very low performance associated with more risky strategies would have a high disutility.

Finally, we can step back and look at design issues. For example, how should p be selected? Small values of p encourage increased variability and reduced correlations, large values of p encourage decreased variability and increased correlations, and values of p near 0.5 do not provide any particular incentive. For contests within organizations, such as explicit or implicit incentive plans in the form of contests, careful consideration should be given to the implications for the organization. For example, if a school wants to be more innovative and nurture high-risk, high-payoff "big ideas," it should decrease p (of tenure) for junior faculty. However, a firm wanting to reduce the chance of very low individual (or aggregate) sales should set $p > 0.5$ in a sales contest. There is little benefit to designing a contest with $p = 0.5$ because it provides no particular incentive in terms of variability.

There are also implications regarding the type of individual who might join the organization. For example, consider a new Ph.D. entering academia with a choice between a school with moderate research expectations and reasonably high p (of tenure) and a top research school with low p but greater rewards associated with winning the tenure contest. An organization wanting to minimize the chance of very low performance and/or to attract people who prefer to stay on well-trodden paths should set p high, whereas an organization wanting to increase the chance of especially high performance (at the cost of an increased chance of especially low performance) and/or to attract people who are competitive and like the challenge of striking off in new directions should set p low.

Appendix

Lemma. *If performance levels are independent, $f_1, \ldots, f_{n-1}$ are symmetric about zero, and $t(x)$ is the pdf of the $(n-k)$th order statistic of the performance levels from $f_1, \ldots, f_{n-1}$, then for all $x > 0$, $t(x) \geqslant t(-x)$ if $p < 0.5$, $t(x) \leqslant t(-x)$ if $p > 0.5$, and $t(x) = t(-x)$ if $p = 0.5$.*

Proof. Consider the case $p < 0.5$, i.e., $k < n-k$. Denote by Ω_i the set of all $n-1$ contestants except C_i, so $|\Omega_i| = n-2$. For each subset d of Ω_i, let $A_d(x) = \prod_{j\in d} F_j(x)$. Then, because all f_j are symmetric about zero, $A_d(-x) = \prod_{j\in d} F_j(-x) = \prod_{j\in d}(1-F_j(x))$. Denote by D_i the collection of all subsets of size $k-1$ of Ω_i. Note that $|D_i| = \binom{n-2}{k-1}$. Let $S_i(d)$ be a one-to-one mapping of D_i to itself such that $\forall\, d$, $d \cap S_i(d) = \varnothing$. For $k-1 = 1$, an example of $S_i(d)$ is $S_i(\{j\}) = \{j+1\}$ if $j < n-2$ and $S_i(\{n-2\}) = \{1\}$. For $k > 2$, the construction of $S_i(d)$ is similar. In this notation, $t(x) = \sum_{i=1}^{n-1} f_i(x) P_i(x)$, where

$$\begin{aligned} P_i(x) &= \sum_{d\in D_i} \prod_{j\in\Omega_i - d} F_j(x) \prod_{j\in d}(1-F_j(x)) \\ &= \sum_{d\in D_i} A_{\Omega_i - d}(x) A_d(-x) \\ &= \sum_{d\in D_i} A_{S_i(d)}(x) A_{\Omega_i - S_i(d) - d}(x) A_d(-x). \end{aligned}$$

Using $S_i^{-1}(d)$, we can write

$$\begin{aligned} P_i(-x) &= \sum_{d\in D_i} A_{\Omega_i - d}(-x) A_d(x) \\ &= \sum_{d\in D} A_{S_i^{-1}(d)}(-x) A_{\Omega_i - S_i^{-1}(d) - d}(-x) A_d(x) \\ &= \sum_{d\in D_i} A_d(-x) A_{\Omega_i - d - S_i(d)}(-x) A_{S_i(d)}(x). \end{aligned}$$

Note that $\Omega_i - S_i(d) - d = \Omega_i - d - S_i(d)$. Therefore, for $x > 0$,

$$\begin{aligned} A_{\Omega_i - d - S_i(d)}(-x) &= \prod_{j\in\Omega_i - d - S_i(d)} F_j(-x) \\ &\leqslant \prod_{j\in\Omega_i - d - S_i(d)} F_j(x) = A_{\Omega_i - S_i(d) - d}(x). \end{aligned}$$

Thus,

$$
\begin{aligned}
P_i(-x) &= \sum_{d\in D_i} A_d(-x) A_{\Omega_i-d-S_i(d)}(-x) A_{S_i(d)}(x) \\
&\leqslant \sum_{d\in D_i} A_d(-x) A_{\Omega_i-S_i(d)-d}(x) A_{S_i(d)}(x) \\
&= P_i(x) \quad \text{for all } i,
\end{aligned}
$$

and because all f_i are symmetric about zero, $t(-x) \leqslant t(x)$. Note that $t(-x) = t(x)$ if and only if for all i, either $f_i(x) = 0$ or $F_j(x) = 0.5$ for all $j \neq i$. The cases $p > 0.5$ and $p = 0.5$ are similar. □

Acknowledgments

The authors thank the associate editor, the referees, and James E. Smith for helpful comments. This research was supported in part by the National Science Foundation under grant SES 98-18855 and INSEAD under R&D Project 2010325. Robert Winkler is grateful to INSEAD, where he was a Visiting Professor at the Singapore and Fontainebleau, France campuses while much of this research was completed. Ilia Tsetlin is grateful to the Fuqua School of Business, Duke University, Durham, North Carolina where he was a Ph.D. student while this research was completed.

References

Brown, K. C., W. V. Harlow, L. T. Starks. 1996. Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *J. Finance* **51**(1) 85–110.

Cabral, L. M. B. 2003. R&D competition when firms choose variance. *J. Econom. Management Strategy* **12**(1) 139–150.

Chevalier, J., G. Ellison. 1997. Risk taking by mutual funds as a response to incentives. *J. Political Econom.* **105**(6) 1167–1200.

Clemen, R. T., R. L. Winkler. 1986. Combining economic forecasts. *J. Bus. Econom. Statist.* **4**(1) 39–46.

Dekel, E., S. Scotchmer. 1999. On the evolution of attitudes toward risk in winner-take-all games. *J. Econom. Theory* **87**(1) 125–143.

Figlewski, S., T. Urich. 1983. Optimal aggregation of money supply forecasts: Accuracy, profitability and market efficiency. *J. Finance* **38**(3) 695–710.

Frank, R. H., P. J. Cook. 1995. *The Winner-Take-All Society*. The Free Press, New York.

Gaba, A., A. Kalra. 1999. Risk behavior in response to quotas and contests. *Marketing Sci.* **18**(3) 417–434.

Hvide, H. K. 2002. Tournament rewards and risk taking. *J. Labor Econom.* **20**(4) 877–898.

Hvide, H. K., E. G. Kristiansen. 2003. Risk taking in selection contests. *Games Econom. Behavior* **42**(1) 172–179.

Lazear, E. P., S. Rosen. 1981. Rank-order tournaments as optimum labor contracts. *J. Political Econom.* **89**(5) 843–864.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2**(1) 71–87.

McCardle, K. F., R. L. Winkler. 1992. Repeated gambles, learning, and risk aversion. *Management Sci.* **38**(6) 807–818.

Moldovanu, B., A. Sela. 2001. The optimal allocation of prizes in contests. *Amer. Econom. Rev.* **91**(3) 542–558.

O'Keeffe, M., W. K. Viscusi, R. J. Zeckhauser. 1984. Economic contests: Comparative reward schemes. *J. Labor Econom.* **2**(1) 27–56.

Pratt, J. W., M. J. Machina. 1997. Increasing risk: Some direct constructions. *J. Risk Uncertainty* **14**(2) 103–127.

Rosen, S. 1986. Prizes and incentives in elimination tournaments. *Amer. Econom. Rev.* **76**(4) 701–716.

Rothschild, M., J. Stiglitz. 1970. Increasing risk: I. A definition. *J. Econom. Theory* **2**(3) 225–243.

Sirri, E. R., P. Tufano. 1998. Costly search and mutual fund flows. *J. Finance* **53**(5) 1589–1622.

Winkler, R. L., D. G. Brooks. 1980. Competitive bidding with dependent value estimates. *Oper. Res.* **28**(3) 603–613.