



Strategic Choice of Variability in Multiround Contests and Contests with Handicaps

ILIA TSETLIN
ANIL GABA
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676

ilia.tsetlin@insead.edu
anil.gaba@insead.edu

ROBERT L. WINKLER*
Fuqua School of Business, Duke University, Durham, NC 27708-0120, USA

rwinkler@mail.duke.edu

Abstract

Variability can be an important strategic variable in a contest. We study optimal strategies involving choice of variability in contests with fixed and probabilistic targets, one-round and multiround contests, contests with and without handicaps, and situations where one contestant can modify variability as well as those in which all contestants have this opportunity. A contestant should maximize variability in a weak position (low mean, high handicap, or low previous performance) and minimize variability in a strong position. In some cases, only these extremes should be used. In other cases, intermediate levels of variability are optimal when the contestant's position is neither too weak nor too strong.

Keywords: contests, targets, variability of performance, handicaps, dynamic strategies

JEL Classification: C44, C61

Introduction

In a contest with a fixed proportion of winners based on relative performance, a contestant should maximize her probability of winning the contest, not her expected performance. Much work on contests has focused on effort level as a key factor (e.g., Lazear and Rosen, 1981; O’Keeffe, Viscusi, and Zeckhauser, 1984; Rosen, 1986; Moldovanu and Sela, 2001). Kalra and Shi (2001) provide an overview of the extensive use of sales contests and investigate incentives for greater effort. Variability can also be an important decision variable. If a contestant can modify the variability of her performance distribution, she can improve her chance of winning. As March (1991, p. 83) states, “In competition to achieve relatively high positions, variability has a positive effect. In competition to avoid relatively low positions, variability has a negative effect.” Inducement for risk taking (in terms of the spread of the performance distribution) in Nash equilibria of one-winner contests is discussed in Dekel and Scotchmer (1999), Hvide (2002), and Hvide and Kristiansen (2003). DeGeorge, Moselle, and Zeckhauser (2004) analyze the risk level chosen by contestants with private information about their quality.

*To whom correspondence should be addressed.

In Gaba, Tsetlin, and Winkler (2004), we develop a model in which the performance of different contestants may be correlated, show that a riskier (less risky) performance distribution is preferred when the proportion of winners in the contest is less (greater) than one-half, where the riskiness is affected by both variability and correlations, and derive an asymptotic formula for the probability of winning. Numerical results with a multinormal model indicate that increasing (decreasing) variability and reducing (increasing) correlations can improve the probability of winning considerably in contests with a low (high) proportion of winners. Moreover, choosing a riskier (less risky) performance distribution is the optimal strategy for a contestant who has the unique opportunity (unavailable to her competitors) to modify riskiness and is also the dominant best-response strategy when all contestants can modify riskiness.

Often a contest has multiple rounds, with new information available at the end of each round. In a sales contest, for example, rewards may be based on total yearly sales. The contestants are able to observe their intermediate sales results during the year and adjust their strategies accordingly. At any given time, a sales representative who is having a particularly good year in terms of sales should not necessarily follow the same strategy as someone else whose sales are lagging far behind their normal pace.

Suppose that a contest consists of a number of rounds, winning is based on total performance (the sum of the performance levels in the individual rounds), and a contestant learns her performance level at the end of each round. Then the optimal strategy to follow in any given round is dynamic, depending on the cumulative performance up to that round and the number of rounds remaining. For example, in match play in golf, a golfer who is three holes ahead with five to play might be expected to play conservatively, whereas the competing golfer who is behind is likely to take chances (e.g., trying a risky shot in an attempt to clear a hazard and reach the green in two shots rather than three). Brown, Harlow, and Starks (1996), Chevalier and Ellison (1997), and Gorjaev, Palomino, and Prat (2001) show that mutual fund managers who are “behind the pack” at mid-year or after three quarters increase the riskiness of their funds relative to funds with better year-to-date returns. Cabral (2003) models R&D races and shows that contestants who are making poor progress should choose a high-risk strategy, taking an approach that has a high variance and is different from the approaches taken by competitors. Contestants who are making good progress should pursue a lower-risk strategy. Gorjaev, Palomino, and Prat (2001) and Cabral (2003) both consider models with two contestants and one winner. Degeorge, Moselle, and Zeckhauser (2004) obtain comparable results in a different setting, showing theoretically and empirically that competitors of lower (higher) quality have more (less) variable performance. Having low (high) quality is comparable to being behind (ahead) in the contest or to having a high (low) handicap, as we will discuss below.

In this article we study a model of a multiround contest and investigate optimal dynamic strategies in a decision-theoretic framework (where only one contestant has the opportunity to modify performance distributions at intermediate stages in the contest, depending on her cumulative performance at those stages) and a game-theoretic framework (where all contestants have this opportunity). A general picture of the optimal multiround strategy can be described as “choose a distribution with large variance if current cumulative performance is low, choose a distribution with medium variance if performance is medium, and choose a

distribution with low variance if performance is high.” Such a strategy differs from a single-round-contest strategy, which will never choose a medium variance when the proportion of winners is different from one-half. Further, the dynamic strategy yields a distribution of total performance that is skewed to the left. The benefits from using the dynamic strategy (compared with having to use the same distribution in each round) are greatest when the proportion of winners is close to one-half.

In many contests, quality differs among contestants. This can be reflected, for example, by differences in mean performance. In our model, we assume that all contestants start out on equal footing, with an initial performance level of zero and mean performance that is identical across contestants. However, because we deal with multiround contests, our model has much wider applicability. Even though the starting point is identical, contestants will of course reach different cumulative performance levels in the intermediate rounds. If we look just at the last round, it can be viewed as a single-round contest with different means. Equivalently, it can be viewed as a single-round contest with different handicaps for the contestants. Here a low mean would be comparable to a high handicap to overcome, and a high mean would be comparable to a low handicap. More generally, any intermediate round can be viewed in the same manner, as a multiround contest with different handicaps or different means. In the mutual fund example noted above, a fund that is doing worse than the competition after three quarters can be viewed as having a handicap going into the last quarter of the year.

Thus, a general picture of the optimal strategy when there are different means or different handicaps can be described as “choose a distribution with large variance if you have a low mean or a high handicap, choose a distribution with medium variance if you have a medium mean or medium handicap, and choose a distribution with low variance if you have a high mean or a low handicap.” This means, for example, that even in a single-round contest, intermediate levels of variability can be optimal. The dividing lines between low, medium, and high means (or high, medium, and low handicaps) will depend on the situation, of course.

Cases with differing means or differing handicaps may exist due to differences in quality among the contestants or simply due to some chance factors. For example, professors in a tenure contest at an academic institution might have different abilities or might start with different initial publication records due to different advisors in graduate school and different luck with early journal submissions. A pharmaceutical firm might have a strong R&D lab or might be fortunate to stumble upon a promising new drug. In a qualifying match for the World Cup, a country’s squad might be less talented than its opponent or the results of earlier qualifying matches might leave it behind the other teams in their group so that it needs to win by several goals in order to qualify. In some cases, handicaps are imposed in a contest in an attempt to make the contest more equitable and more interesting for contestants of lower quality. Larkey et al. (1997, p. 606), in their paper on skill in games, note that “in a few games such as golf there is elaborate ‘handicapping’ to create fair contests among unequally skilled players.”

Since the key issue in our multiround contest model is the optimal strategy for a contestant in the intermediate rounds, when some contestants are ahead and others behind, it is useful to build some intuition by first discussing a model where contestants with different handicap

levels face a fixed target (rather than a rank-order contest which implies a probabilistic target) in a single round. We do so in Section 1. Our model of a multiround contest is developed in Section 2. Numerical results for the multiround contest are given in Section 3.1 for the decision-theoretic framework and in Section 3.2 for the game-theoretic framework. Section 4 presents a brief summary and discussion of the results for multiround contests and their implications for contests with different means or handicaps.

1. A contest with a fixed target

Contests with fixed targets are quite common. A prime example of fixed targets that are explicitly set is the use of quotas for salesforce compensation (e.g., Churchill, Ford, and Walker, 1993). An instructor may set fixed targets for grades, such as a requirement of an average of 90 to earn an A. Fixed targets may be perceived even if they are not explicitly given. For example, firms may engage in earnings management to reach targets such as positive profits, profits greater than the preceding period, and profits in excess of analysts' expectations (DeGeorge, Patel, and Zeckhauser, 1999).

Consider a fixed-target or quota setting where a contestant receives payoff W if her performance is at least as high as a target t and gets payoff $L < W$ if her performance is below t . Denote her expected performance by y . Equivalently, we can think of the contestant having an initial starting point y and subsequent performance x with $E(x) = 0$, where the payoff W is received if the final performance level $y + x \geq t$. Then $h = t - y$ can be thought of as a handicap for the contestant, with a positive (negative) handicap implying that the contestant is starting below (above) the target. While y can be thought of as a mean or as a starting point that imposes a handicap, we will use the handicap interpretation in this section.

Suppose that the contestant can choose the variability of her performance and that there is no cost associated with this choice. Denote the standard deviation of her performance by $s > 0$ and suppose that x/s is symmetrically distributed about zero with cdf F . The contestant should choose s to maximize her probability of attaining a performance x at least as high as $t - y$ so that $y + x \geq t$. This probability is given by

$$\begin{aligned} P(\text{win} | t, y, s) &= P(y + x \geq t | t, y, s) = 1 - F[(t - y)/s] = 1 - F(h/s) \\ &= P(\text{win} | h, s). \end{aligned}$$

From Gaba and Kalra (1999), $P(\text{win} | h, s)$ is maximized at the highest (lowest) available s for a positive (negative) handicap level h . It is rather intuitive that if the handicap is negative (i.e., the contestant's expected final performance level is above the target), then it is optimal to choose the lowest available s . Similarly, if the handicap is positive, so that the contestant's expected final performance level is below the target, it is best to maximize s . In other words, the optimal s does not depend on the value of h , but only on its sign. However, we would expect that the benefits from manipulating s would depend upon the magnitude of h . For example, given $h > 0$, a contestant with a higher h would benefit more for a given increase in s ; the more one is below the target the greater the benefit from a given increase in variability

of performance. Similarly, for $h < 0$, lower values of h would suggest greater benefit from a given decrease in performance variability.

One way to illustrate how the benefits of changing s are related to the handicap levels is to consider the maximum penalty, in terms of an increase in h , that a contestant would be willing to incur for a given change in s . For positive handicap levels, a contestant with a higher h would be willing to pay a higher penalty for a given increase in s . Similarly, for negative handicap levels, a contestant with a lower h would be willing to pay more for a given reduction in s . The proposition below builds upon such an idea.

Proposition 1. (a) All points (h, s) satisfying $h = sF^{-1}(1 - p)$ have the same probability p of winning. (b) Let $s_1 > 0, s_2 > 0$. A contestant prefers (h_2, s_2) to (h_1, s_1) if $h_2/s_2 < h_1/s_1$ and is indifferent between (h_1, s_1) and (h_2, s_2) if $h_2/s_2 = h_1/s_1$. The preference is strict if $0 < F(h_1/s_1) < 1$ and $F(z)$ is strictly increasing for all z such that $0 < F(z) < 1$.

Proof: For (a), $h = sF^{-1}(1 - p)$ is the solution of $P(\text{win} | h, s) = p = 1 - F(h/s)$. For (b), $P(\text{win} | h, s)$ is nonincreasing in h/s since F , as a cdf, is nondecreasing, and $P(\text{win} | h, s)$ is strictly decreasing if F is strictly increasing. \square

From Proposition 1, the indifference curves for h and $s > 0$, given $0 < p < 1$, are linear, of the form $h = sF^{-1}(1 - p)$. The probability of winning is p along this line, and if F is strictly increasing the probability of winning is strictly higher (lower) than p for all points below (above) this line. Figure 1 illustrates these indifference lines for the case where performance is normally distributed (i.e., $F(x) = \Phi(x)$ is the standard normal cdf).

When $p < 0.5$, the slope $F^{-1}(1 - p)$ of an indifference line equals the maximum increase in handicap that a contestant is willing to adopt for changing the standard deviation from s to $s + 1$. For $p > 0.5$, the slope is negative and represents the minimum decrease in handicap that a contestant needs as compensation for going from s to $s + 1$. Alternatively,

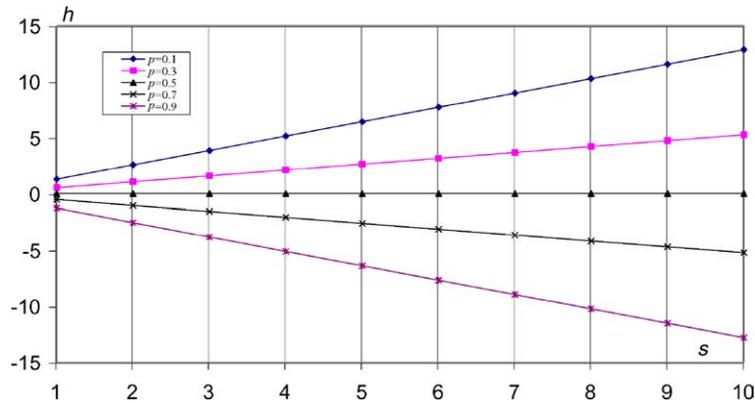


Figure 1. Indifference lines for (handicap, standard deviation) pairs with a normal distribution of performance. A contestant has the same probability p of winning for any two points that are on the same line.

it can be interpreted as the maximum increase in handicap that a contestant will accept for going from s to $s - 1$. Since $F^{-1}(1 - p)$ is decreasing in p and $F^{-1}(0.5) = 0$, these changes in handicaps to maintain indifference are larger as p gets closer to zero or one. For example, if the probability of winning in Figure 1 is 0.1 (0.3), the contestant is willing to increase her handicap by 1.28 (0.52) in order to increase s by 1. Similarly, if her probability of winning is 0.9 (0.7), she is willing to increase her handicap level by 1.28 (0.52) in order to decrease s by 1.

To sum up, a contestant facing a fixed target will simply adopt the maximum or minimum variability depending on whether her handicap is positive or negative. In this setting an intermediate level of variability should be considered only with a handicap of zero (implying $p = 0.5$), in which case the probability of winning is the same for any level of variability. A contestant having a more severe positive (negative) handicap is willing to pay more for a given increase (decrease) in the variability of performance. With these intuitively reasonable results in mind, we turn to the case of a multiround contest with a probabilistic target instead of a fixed target.

2. A model of a multiround contest with a probabilistic target

Assume that a contest has $n \geq 2$ contestants (C_1, \dots, C_n), and consists of $m \geq 1$ rounds. The performance levels of all n contestants are independent across contestants and across rounds. The contestants with the k highest values of total performance (the sum of the performance levels for the m rounds), $1 \leq k \leq n - 1$, are winners in the contest, each receiving payoff W . Each of the remaining contestants receives payoff $L < W$.

The proportion of winners in the contest, k/n , is denoted by p . Because the proportion of winners is fixed, whether a contestant wins or not depends not just on her absolute performance but on her relative performance compared to the other contestants. Further, since the performance of her rivals is uncertain, she faces a probabilistic target.

We consider two frameworks: a decision-theoretic framework and a game-theoretic framework. In the decision-theoretic framework, only one of the contestants, C_n , can modify the distribution of her performance level. We assume that C_n wants to maximize her probability of winning and that she incurs no cost in choosing distributions. The set of available distributions is denoted by Ψ . Before each round, C_n chooses her distribution $f \in \Psi$ for the performance in that round and then observes the resulting performance. All other contestants use a fixed distribution (known to contestant C_n) for their performance levels in each round. The dynamic component of the problem arises because C_n can choose a distribution in a particular round as a function of her current cumulative performance after the previous round. In the game-theoretic framework, *each* contestant chooses a distribution in each round as a function of her current cumulative performance. In other words, the dynamic choice of performance distributions is extended to all contestants.

We first describe the general solution of the dynamic problem in the decision-theoretic framework, and then extend it to the case of the game-theoretic framework. In the decision-theoretic framework, we find the distribution $f \in \Psi$ that is optimal for C_n in round j if her cumulative performance after round $j - 1$ is x_{j-1} , assuming that optimal strategies are chosen in all subsequent rounds and setting $x_0 = 0$ as the starting point. Denote the

resulting probability of winning, found at the beginning of round j as a function of x_{j-1} , by $P_j(x_{j-1})$. Note that $P_1(0)$ is the probability of winning at the beginning of the contest.

Proposition 2. *Let $T(x_m)$ represent the probability that C_n wins given that her total performance is x_m . Since 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 $n - 1$ total performance levels of the other contestants. Then $P_j(x_{j-1})$ is recursively found as follows:*

$$P_m(x_{m-1}) = \max_{f \in \Psi} \int_{-\infty}^{+\infty} f(u)T(x_{m-1} + u) du$$

and

$$P_j(x_{j-1}) = \max_{f \in \Psi} \int_{-\infty}^{+\infty} f(u)P_{j+1}(x_{j-1} + u) du$$

for $j = m - 1, \dots, 1$.

This dynamic program does not have an analytical solution when $m > 1$. We study the problem numerically, using the following model. Let $\Psi = \{f_s \mid s = 1(0.1)9\}$, where f_s is the pdf of the normal distribution with mean zero and standard deviation s . Thus, s ranges from 1 to 9, with discrete steps of 0.1 small enough to approximate the situation where any s between 1 and 9 is available and a ratio $s_{\max}/s_{\min} = 9$ large enough to provide a reasonable range of values. This is equivalent to any setting with $s_{\max} = 9s_{\min}$.

Solving the dynamic program requires numerical integration and requires calculating optimal strategies for any feasible scores x_{j-1} , $j = 1, \dots, m$. To facilitate the solution process, we discretize performance levels with steps of size 0.2 and we truncate the normal distribution f_s at $\pm 3s$, normalizing the truncated distribution. The step of 0.2 is small enough for an adequate approximation of the normal distribution yet large enough for fast numerical calculations. Note that in the discrete case there is a problem of resolving ties, which is done as follows: If q contestants have the same total performance $x_m = x$ and there are a contestants with total performance above x , where $k - q < a < k$, then all q contestants with total performance x have equal chances of winning, namely $(k - a)/q$.

In the decision-theoretic setting, the dynamic program is a discrete numerical problem that can be solved by backward induction in a finite number of steps. In the game-theoretic setting, we iteratively solve for a symmetric Nash equilibrium. In the first step we assume that all contestants except C_n use some s in each round and we find the best reply for contestant C_n . In the next step we find the best reply for contestant C_n if all rivals use C_n 's optimal strategy from the previous step. The iterative search ends as soon as the optimal strategies coincide for two consecutive steps, and the result is a symmetric Nash equilibrium since it is the best reply assuming that all rivals use the same strategy. We do not make any claims about whether this equilibrium exists for all possible sets of model parameters, although it exists in all cases that we analyzed. Furthermore, in all cases that we have considered the symmetric equilibrium strategy is the same for many different

choices of an initial distribution for the first iteration step. Thus, we feel that the results presented in Section 3.2 describe the unique symmetric equilibrium for each set of model parameters.

An important characteristic of our multiround contest is that although each contestant has a starting point of zero, so that no one is handicapped relative to the competition at the beginning of the contest, the contestants will have de facto handicaps after one or more rounds due to differing performance. In particular, the last round is itself equivalent to a single-round contest with handicaps, differing from the contest in Section 1 because the target is probabilistic rather than fixed. Therefore, the model has implications not just for multiround contests, but also for single-round probabilistic-target contests with handicaps.

3. Numerical results for the multiround contest

We consider examples with $n = 10$ contestants, the proportion of winners $p = 0.1(0.2)0.9$, and the number of rounds $m = 1(1)4$. As noted in Section 2, the set of available distributions is $\Psi = \{f_s \mid s = 1(0.1)9\}$, where f_s is the pdf of the normal distribution with mean zero and standard deviation s . Thus, a distribution is specified by the choice of s , with the minimum at 1 and the maximum at 9.

3.1. Decision-theoretic framework

Figure 2 shows the optimal strategies for C_n in rounds 2, 3, and 4 as a function of current cumulative performance x in a contest with $m = 4$ rounds for the case $p = 0.3$. In each round, all rivals use a fixed- s strategy with $s = 9$, the maximum variability, which is also the optimal strategy in a one-round contest. (Note that if round-by-round performance is

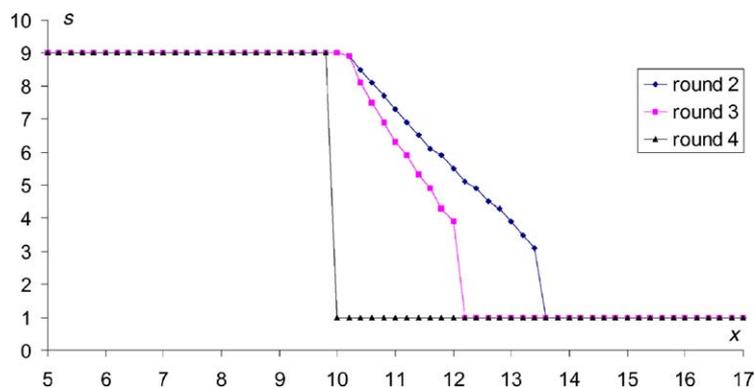


Figure 2. Optimal choice of s for C_n in rounds 2, 3, and 4 as a function of cumulative performance x in a contest with $m = 4$ rounds and proportion of winners $p = 0.3$, when other contestants use $s = 9$ in each round.

unobservable, then a multiround contest is strategically equivalent to a one-round contest.) In the first round, $x = 0$ and the optimal strategy for C_n is to set $s = 9$. In the last round it is always optimal to use either the largest or smallest variance. However, in the preceding rounds there is a range of values of x (approximately from 10.2 to 13.4) for which some medium variance is optimal. This demonstrates that in a multiround contest, unlike a single-round contest, intermediate strategies are sometimes optimal. With $p < 0.5$, C_n should start at s_{\max} in the first round and then reduce s in later rounds if x becomes high enough. Similarly, if $p > 0.5$, C_n should start at s_{\min} , increasing s in later rounds if x becomes low enough. This is consistent with the idea that a contestant who is doing well should follow a lower-risk strategy. In each round, as x gets even higher the optimal value of s decreases until it eventually reaches $s_{\min} = 1$, at which point the contestant feels comfortable enough about winning to want to preserve her current performance level.

As noted in Section 2, the last round is equivalent to a one-round contest with a probabilistic target in which contestants have different handicaps. In the last round it is optimal for C_n to use the largest or the smallest variance, as in the case of a fixed target discussed in Section 1. However, in the intermediate rounds it is sometimes optimal for C_n to use intermediate variances, depending on her cumulative performance.

Figure 3 presents the increase in C_n 's probability of winning due to using the optimal dynamic strategy instead of the optimal fixed- s strategy (using the s that is optimal when performance levels in intermediate rounds are unobservable) when the rivals use the optimal fixed- s strategy ($s = 9$ when $p \leq 0.5$ and $s = 1$ when $p \geq 0.5$). These benefits are shown for $m = 2(1)4$ as a function of p . Note that when $m = 1$, there is no benefit since the contest has only one round and every contestant uses the optimal strategy; each contestant's probability of winning is p . As expected, the benefits of using the dynamic optimal strategy increase with m , since a larger m gives C_n more opportunities to change the distribution contingent on her current performance. The improvements are also greater for values of p closer to 0.5 since for low p large s is optimal with high probability, and for high p low

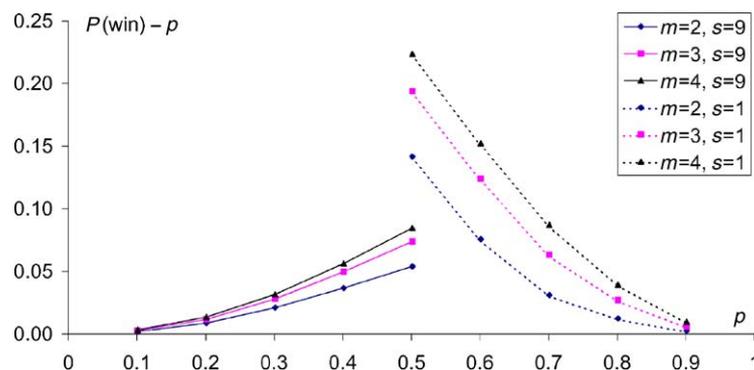


Figure 3. Improvement in the probability of winning for C_n , $P(\text{win}) - p$, from optimal dynamic strategies for $p = 0.1(0.1)0.9$ in a contest with $m = 1(1)4$ rounds when other contestants use $s = 9$ for $p \leq 0.5$ and $s = 1$ for $p \geq 0.5$ in each round.

s is “almost always” optimal. Thus, as p gets close to zero or one, the optimal dynamic strategy seldom strays from the optimal fixed- s strategy. Finally, the benefits are higher for a given $p > 0.5$ than for the corresponding $1 - p$; the dynamic strategy can improve the chances of winning in a contest with many winners more effectively than in a contest with few winners.

3.2. Game-theoretic framework

We now consider the case where all contestants observe their own performance levels after each round and then choose s for the next round depending upon their cumulative performance levels. In all the cases we have investigated, a symmetric Nash equilibrium exists. By symmetry, the probability of winning is p for each contestant under such an equilibrium. Here following the indicated strategy is necessary to avoid reducing the probability of winning below p . This is in contrast to the decision-theoretic framework, where following the optimal strategy yields $P(\text{win}) \geq p$ because the rivals do not have the opportunity to modify their distributions dynamically.

Figure 4 presents the optimal equilibrium strategy for each contestant in rounds 2, 3, and 4 as a function of current cumulative performance x in a contest with $m = 4$ rounds and $p = 0.3$. As in Figure 2, s_{\max} is optimal for low values of x , s_{\min} is optimal for high values of x , and there is a range of values of x for which intermediate values of s are optimal. This range of values is roughly from 12.5 to 18.5, as compared with 10.2 to 13.4 in Figure 2. When all contestants are following the optimal dynamic strategy, the anticipated total performance required to win (i.e., the probabilistic target) is much higher and shifts from s_{\max} therefore occur at a higher cumulative performance x . Furthermore, the optimal strategies in the game-theoretic framework do not differ as much from round to round as they do in the decision-theoretic framework. In particular, intermediate strategies are just as

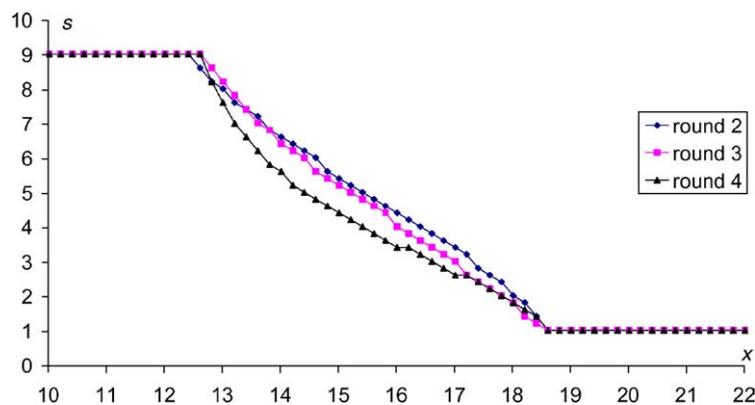


Figure 4. Optimal choice of s for C_n in rounds 2, 3, and 4 as a function of cumulative performance x in a contest with $m = 4$ rounds and proportion of winners $p = 0.3$, when all contestants follow symmetric Nash equilibrium strategies.

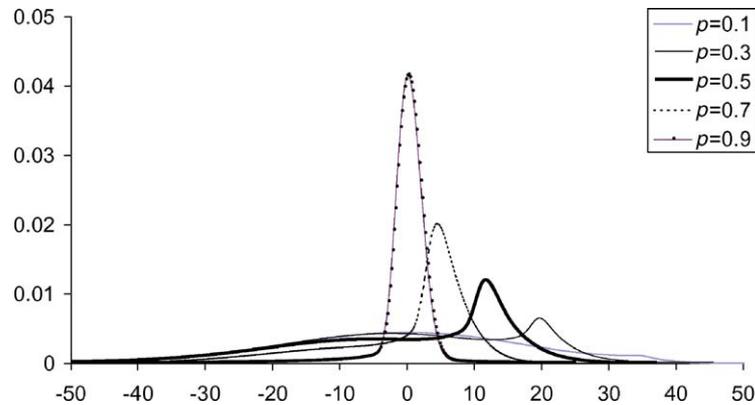


Figure 5. Probability density functions of total performance for each contestant under the symmetric Nash equilibrium for $p = 0.1(0.2)0.9$ in a contest with $m = 4$ rounds.

prevalent in the last round as in rounds 2 and 3, in contrast with Figure 2 where intermediate strategies are never optimal in the last round. Since the last round is equivalent to a single-round contest with handicaps, these results show that intermediate strategies can be optimal in a single-round contest.

Figure 5 shows the pdfs of total performance for the symmetric Nash equilibrium strategy for $p = 0.1(0.2)0.9$ and $m = 4$. For $p = 0.9$ the distribution of total performance almost coincides with a normal pdf with $s = 1$, since the contestants seldom stray from $s = 1$ in this case. However, as p decreases to 0.7, 0.5, and 0.3, the distributions become skewed to the left with modes at 4.4, 11.6, and 19.6, respectively. There is a second, smaller mode to the left at -2.3 for $p = 0.3$ and at -4.8 for $p = 0.5$ (they are tiny “bumps,” not visible in the graph). When cumulative performance is low, a high- s strategy is optimal, increasing the spread of the lower half of the distribution. As p decreases, the right-hand mode shifts to the right and gets lower. For example, the mode of 19.6 when $p = 0.3$ is slightly higher than 18.5, the value of x beyond which the optimal s is s_{\min} in Figure 5. Evidently a contestant with $x > 18.5$ should feel comfortable enough about winning to minimize her standard deviation, and the result is a concentration of probability in the region just above 18.5. When $p = 0.1$, the main mode is at zero and the right-hand mode is a relatively small bump on the pdf at 34.2. For such a low value of p , the contestants will only depart from $s = 9$ if they are fortunate enough to reach a cumulative performance that is quite high. The probability of doing so is low, which is why the overall performance distribution when $p = 0.1$ is close to a normal pdf with $s = 9$ and the right-hand mode is just a small bump.

As expected, the standard deviation of total performance decreases as p increases in all the cases that we considered for $m = 4$; it is 17.67 for $p = 0.1$, 16.87 for $p = 0.3$, 15.38 for $p = 0.5$, 10.96 for $p = 0.7$, and 3.14 for $p = 0.9$. We conjecture that in general, for $m > 1$, the equilibrium standard deviation of total performance decreases as p increases, although we are unable to prove this. In contrast, if round-by-round performance is unobservable and a fixed- s strategy is therefore optimal, the decrease in standard deviation

of total performance is a single step, from the largest possible standard deviation (17.77) for $p < 0.5$ to the smallest possible standard deviation (1.98) for $p > 0.5$. Note that these standard deviations are slightly different from the corresponding theoretical values. Choosing $s = 9$ in all $m = 4$ rounds should yield a standard deviation of $s(m^{1/2}) = 18$, and choosing $s = 1$ should yield a standard deviation of 2. Due to the truncation of the normal distribution at $\pm 3s$, the corresponding standard deviations are 17.77 and 1.98.

Allowing the contestants to observe their performance levels after each round and follow a dynamic strategy decreases (increases) the standard deviation of total performance for $p < (>) 0.5$. Also, the standard deviation of total performance is greater in the game-theoretic setting than in the decision-theoretic setting. A comparison of Figures 2 and 4 shows that a larger cumulative performance level is required to cause s to decrease from s_{\max} in the game-theoretic setting, reflecting the fact that the probabilistic target is higher. Thus, for a range of cumulative performance levels, a larger s is optimal in the game-theoretic setting, so total performance is more variable.

4. Summary and discussion

The primary thrust of this paper is to focus on the choice of variability as an important instrument for a contestant and to investigate how optimal strategies involving choice of variability differ from situation to situation. We have considered a variety of settings that are common in real-world contests. In particular, our results address:

- contests with fixed targets and contests with probabilistic targets;
- one-round contests and multiround contests;
- contests where all contestants are on equal footing and contests where there are differences in quality (mean performance) among contestants or contestants have different handicaps;
- situations where one contestant has a unique opportunity to change variability from round to round and situations where all contestants have this opportunity.

In all of these settings, the optimal strategy for a contestant is to choose the highest possible variability when she is in a weak position (due, for example, to a low mean performance, a low cumulative performance at an intermediate point in a multiround contest, or a high handicap) and to choose the lowest possible variability when she is in a strong position (a high mean, a high cumulative performance at an intermediate point, or a low handicap). In some cases, only these extreme levels of variability should be used, and there is a cutoff point (in terms of the handicap, the mean, or the cumulative performance) at which a contestant would switch from the highest possible variability to the lowest possible variability. In other cases, however, intermediate levels of variability can be optimal when a contestant's position is neither too weak nor too strong.

In a one-round fixed-target setting, a contestant with a positive (negative) handicap should always choose the maximum (minimum) variability. Moreover, a contestant with a more severe positive (negative) handicap should be willing to pay more for a given increase (decrease) in variability. For example, in a quota-based compensation scheme a salesperson who is further below the quota should be willing to pay more to increase variability. The

salesperson might be more willing to seek out new prospective customers for whom the probability of getting an order is small but the size of the order is large.

In the model of a multiround contest presented in Section 2, the target is probabilistic because the proportion of winners is fixed and each contestant is uncertain about the performance of the other contestants. The contestants start out with no handicaps, but after one or more rounds they will effectively have handicaps because of differing performance that causes some to be ahead and some to be behind. Thus, at any point after the first round, the contest can be viewed as one with handicaps and a probabilistic target. Moreover, although the model has equal means for the contestants, the differences in cumulative performance at any point are equivalent to having identical cumulative performance levels but different means (and/or different handicaps).

In a decision-theoretic framework for a multiround contest, only one contestant has the opportunity to observe her performance in intermediate rounds and to modify her variability in an optimal dynamic manner. The other contestants cannot see their intermediate performance levels and therefore will follow the best fixed-variability strategy, which is to use the maximum (minimum) variability if the proportion of winners is less than (greater than) one-half. For the contestant who can modify variability, it is always optimal to use the largest or smallest variability in the last round. The last round is equivalent to a one-round contest where contestants have different handicap levels, and these results for a situation in which contestants are chasing a probabilistic target are consistent with the one-round fixed-target setting.

At an intermediate stage in the multiround contest (i.e., a round other than the first or last round), the contestant who can observe her performance should choose the maximum variability if her cumulative performance is low, a medium variability if her cumulative performance is medium, and the minimum variability if her cumulative performance is high. A major difference between one-round and multiround contests is that standard deviations other than s_{\min} and s_{\max} are sometimes optimal in multiround contests but never in one-round contests (except perhaps when $p = 0.5$, where the probability of winning is 0.5 for any s). The high/low levels of current performance between which intermediate standard deviations should be used depend upon the situation, particularly upon the proportion of winners p and the number of rounds remaining. These high/low levels increase as the proportion of winners p decreases, and the spread between the high and low levels increases with the number of remaining rounds.

The improvement in the probability of winning due to using the optimal dynamic strategy instead of the optimal fixed-variability strategy increases with m , the number of rounds. As m increases, there are more opportunities to take advantage of the dynamic strategy. Also, the improvement is greatest when the rivals use a low s and when p is near 0.5. If p is low, then there is only a small probability of having an intermediate cumulative performance level that is high enough to warrant using a low s . Thus, choosing s_{\max} in each round does not hurt much, compared to the optimal dynamic strategy. Similarly, for large p there is a small probability of having a low enough intermediate cumulative performance level to use a high s , so that using s_{\min} in all rounds does not hurt much. If p is closer to 0.5, however, then the benefits of choosing the standard deviation as a function of the current performance are higher.

Note that in the decision-theoretic setting, the source of the uncertainty about the target is not important. The uncertainty about the target in our model comes from uncertainty about competitors' performance, but it could equally well arise in other ways (e.g., in a quota-based compensation scheme where a contestant is not sure what the quota is). Even if winning the multiround contest does not depend on how other contestants fare, the same recursive formulas from Proposition 2 apply with T representing the cdf of the uncertain target.

In the game-theoretic framework of the multiround contest, where all contestants have the opportunity to observe their own performances and modify their variability in intermediate rounds, standard deviations other than s_{\min} and s_{\max} are sometimes optimal not just in the intermediate rounds but also in the last round. Since the last round can be viewed as a single-round contest with handicaps and a probabilistic target, this demonstrates that intermediate levels of variability can be optimal for some handicaps in a single-round contest with a probabilistic target. Further, the high/low levels of current performance between which standard deviations other than s_{\min} and s_{\max} are optimal are higher compared to those in the decision-theoretic framework. This makes sense because all contestants are able to benefit from observing their performance and making round-to-round adjustments in variability, thereby improving their expected lot and making the contest more competitive.

The dynamic strategies yield a distribution of total performance for each contestant that is skewed to the left, with the skewness being most pronounced at intermediate values of p . For a given m , the standard deviation of total performance increases as p decreases. Note that since all distributions in each round are symmetric about zero, the expected value of the total performance is also zero, but the median and mode shift to the right and the left tail becomes fatter.

In our analysis, we assume that there are no costs associated with modifying variability. Intuitively, one might think of the increased risk of a very low performance as a cost. However, as long as we have the "bang-bang" reward structure in which the only thing that matters is reaching the target or not, an absolutely awful performance is no worse than a "pretty good" performance that falls just short of the target. For example, in the last "real" stage of the 2003 Tour de France, with just the ceremonial ride into Paris to follow, the second-place cyclist crashed as a result of taking chances on a wet surface in an attempt to catch up to the leader. The crash resulted in a bad performance for that stage, eliminating any chance of winning the Tour de France, but otherwise left the cyclist no worse off.

Imposing a cost on modifications in variability would modify the reward structure and possibly lead to greater use of intermediate levels of variability. In the case of a fixed target, we observe that a contestant who is closer to the target in terms of expected performance (i.e., a contestant with a lower absolute handicap) is willing to pay less for a given change in the standard deviation, as per Proposition 1. If costs for modifying variability do exist, with increasing marginal costs for larger changes in variability, we would expect that a contestant with a low absolute handicap might be inclined to adopt some intermediate rather than extreme modification in variability. Of course, the optimal strategy would depend on the exact shape of the cost function. Similarly, in a multiround contest we would anticipate that contestants would be willing to pay less for any given change in variability and that this would expand the range of values of cumulative performance for which intermediate values are optimal. Other than this, the overall results are unlikely to change.

Our model also assumes that a contestant who has the opportunity to observe her own performance level each round does not have information about the performance levels of the other contestants. In many cases, this assumption is realistic (e.g., sales contests with sales representatives in different districts). In other cases, it is a good approximation because information about the performance of rivals might be quite noisy or might be available only after a significant time lag. Relaxing the assumption makes the problem much more difficult to solve, especially as the number of contestants increases, because it is necessary to find optimal strategies for all possible vectors (x_1, \dots, x_n) of cumulative performance levels of the n contestants at each round. In any event, information about rivals' performance should not change the general nature of the results, which are based on expectations about rivals' performance. If the rivals are doing better (worse) than expected, then the high/low levels of current performance (or, correspondingly, low/high handicaps) between which intermediate standard deviations should be used would likely shift up (down) a bit, making the use of larger (smaller) standard deviations more likely. Similarly, although numerical results would change, the general picture of the optimal strategy and the skewed shape of the distribution of total performance should be robust with respect to non-normal performance distributions and different values for the numbers of contestants and rounds.

We have focused on the strategies used by contestants given particular details of the contest. There are also many interesting issues on the design side of the coin. For example, a contest designer might want to give each contestant a "fair" chance to win in order to make the contest more interesting both for contestants and observers and to maintain incentives. To this end, handicaps are often used (e.g., in golf). In multi-round contests, interest can be maintained by making early rounds count less; this is observed in some betting pools for basketball tournaments. On the other hand, there are contests in which the leaders are given some advantages or past performance is rewarded. Tournaments in sports are often seeded so that the contestants coming into the tournament with strong records are rewarded by playing early tournament games against contestants with weak records. This is done in tennis and in many of the major team sports. In other cases, such as skating and golf, the leading contestant going into the last round is scheduled to go last, enabling her to observe how the other contestants have done. This provides incentive for those not in the lead to take risks in an attempt to catch up to the leader. Then, if a challenger posts an exceptionally good score, the leader will see this and may have to take some chances (increasing variability) to stay in the lead. If the challengers do not do so well, the leader can adopt a conservative strategy (decreasing variability). In the Tour de France example mentioned above, the leader was behind his rival and slowed down upon learning that the rival had crashed. Sequential play changes the incentives somewhat compared to simultaneous play. These design issues tend to deal with proper incentives, equity, rewards for good performance, and maintaining interest in the contest.

Acknowledgment

This research was supported in part by the National Science Foundation under grant SES 98-18855 and by the Center for Decision Making and Risk Analysis at INSEAD. We thank Richard Zeckhauser for helpful comments that led us to broaden the scope of the paper.

Winkler is grateful to INSEAD, where he was a Visiting Professor at the Singapore and Fontainebleau campuses while much of this research was completed. Tsetlin is grateful to the Fuqua School of Business, Duke University, where he was a Ph.D. student while much of this research was completed.

References

- Brown, K.C., W.V. Harlow, and L.T. Starks. (1996). "Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry," *Journal of Finance* 51, 85–110.
- Cabral, L.M.B. (2003). "R&D Competition When Firms Choose Variance," *Journal of Economics and Management Strategy* 12, 139–150.
- Chevalier, J. and G. Ellison. (1997). "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy* 105, 1167–1200.
- Churchill, G.A., N.M. Ford, and O. Walker. (1993). *Sales Force Management: Planning, Implementation and Control*. Homewood, IL: Richard D. Irwin.
- DeGeorge, F., B. Moselle, and R. Zeckhauser. (2004). "The Ecology of Risk Taking," *Journal of Risk and Uncertainty* 28, 195–215.
- DeGeorge, F., J. Patel, and R. Zeckhauser. (1999). "Earnings Management to Exceed Thresholds," *Journal of Business* 72, 1–33.
- Dekel, E. and S. Scotchmer. (1999). "On the Evolution of Attitudes Toward Risk in Winner-Take-All Games," *Journal of Economic Theory* 87, 125–143.
- Gaba, A. and A. Kalra. (1999). "Risk Behavior in Response to Quotas and Contests," *Marketing Science* 18, 417–434.
- Gaba, A., I. Tsetlin, and R.L. Winkler. (2004). "Modifying Variability and Correlations in Winner-Take-All Contests," *Operations Research* 52, 384–395.
- Goriaev, A., F. Palomino, and A. Prat. (2001). "Mutual Fund Tournament: Risk Taking Incentives Induced by Ranking Objectives," CEPR, Tilburg University, working paper.
- Hvide, H.K. (2002). "Tournament Rewards and Risk Taking," *Journal of Labor Economics* 20, 877–898.
- Hvide, H.K. and E.G. Kristiansen. (2003). "Risk Taking in Selection Contests," *Games and Economic Behavior* 42, 172–179.
- Kalra, A. and M. Shi. (2001). "Designing Optimal Sales Contests: A Theoretical Perspective," *Marketing Science* 20, 170–193.
- Larkey, P., J.B. Kadane, R. Austin, and S. Zamir. (1997). "Skill in Games," *Management Science* 43, 596–609.
- Lazear, E.P. and S. Rosen. (1981). "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy* 89, 843–864.
- March, J.G. (1991). "Exploration and Exploitation in Organizational Learning," *Organization Science* 2, 71–87.
- Moldovanu, B. and A. Sela. (2001). "The Optimal Allocation of Prizes in Contests," *American Economic Review* 91, 542–558.
- O'Keefe, M., W.K. Viscusi, and R.J. Zeckhauser. (1984). "Economic Contests: Comparative Reward Schemes," *Journal of Labor Economics* 2, 27–56.
- Rosen, S. (1986). "Prizes and Incentives in Elimination Tournaments," *American Economic Review* 76, 701–716.