



AMERICAN UNIVERSITY
W A S H I N G T O N , D C

Department of Economics
Working Paper Series

**Adapting Strategic Risk in Corporate
Tournaments**

By:

James Bono

No. 2008-17

Fall 2008

<http://www.american.edu/cas/econ/workpap.htm>

Copyright © 2008 by James Bono. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

Adapting Strategic Risk in Corporate Tournaments

James W. Bono
Economics Department
American University

October 1, 2008

Abstract

The way in which agents manipulate the distribution of performance outcomes in strategic settings has received increasing attention in the game theory literature. This paper uses an evolutionary approach to examine the optimal adaptation of strategic variability in corporate promotion tournaments. The model describes a situation in which agents are promoted to a higher salary based on observable performance, which depends stochastically on effort. Simulation results show that the optimal adaptation of risk-taking is highly dependent on the population mix. However, strategies that involve adapting risk early in the tournament are almost never part of an evolutionarily stable state, particularly when using uniform initial conditions. Results also show how managers can choose rank-order payoff schemes and tournament lengths to optimize with respect to risk-taking and effort.

JEL Classification Codes: C7, C6, D8, M5

KEYWORDS: Evolutionary Game Theory, High Variance Strategies, Tournaments.

† I would like to thank Professors Donald Saari and Michael McBride, the Economics Department and the Institute of Mathematical Behavioral Sciences at U.C. Irvine as well as two anonymous referees.

1 Introduction

Many strategic environments provide players with the opportunity to strategically manipulate the variance of their performance. In their famous paper, Brown, Harlow, and Starks (1996), show that mutual fund managers utilize the riskiness of their portfolio as a strategic variable in competition with other funds. For R&D competitions, Cabral (2003) demonstrates how firms similarly form strategies that involves the variability of research activities. Also, the strategic choice of variance in sports has been examined by Bronars and Oettinger (2001) with respect to golf tournaments and Grund and Gürtler (2005) with respect to soccer matches.

This paper introduces corporate promotion contests as a new setting for the study of how players strategically manipulate variance in games. Following the economic literature on corporate hierarchies as tournaments [see Bognanno (2001)] this study focuses on how the tournament design affects the equilibrium level of risk-taking and contrasts big risk and big reward strategies with small risk and small reward strategies. Among the key results are the findings that the level of risk-taking can be optimized with respect to the length of the tournament (number of promotion levels) and the convexity of the payoff structure. Results also show that while the strategic interactions are rich enough that there are often no dominant strategies, certain strategies are especially unlikely in equilibria. In particular, strategies that adapt from high variance to low variance when ahead before the half-way point are very unlikely to exist in equilibrium.

One of the main foci of the literature is how agents adjust their risk-taking depending on their relative position at a given point in the tournament. In addition to the work of Brown et al. (1996), Chevalier and Ellison (1997) and Gorjaev, Palomino, and Prat (2001) provide empirical evidence that mutual fund managers increase the risk of their funds when earnings are low relative to other funds and decrease risk when earnings are high. Gorjaev et al. (2001) also construct a two-period, two agent model to show that such a strategy is part of a Nash equilibrium. Cabral (2003) shows that in R&D races, competitors that lag behind should pursue high risk strategies, while competitors that are ahead should “play it safe.” Bronars and Oettinger (2001) analyze data from professional golf tournaments and find that players that are trailing on the final holes shoot highly atypical scores because they are attempting higher risk shots.

Though all of this literature identifies the existence of strategies that adapt variance depending on tournament position, it is not at all clear how this adaptation is done optimally. This is especially true when it comes to the tournament payoff scheme. Yumoto (2003) and Gaba, Tsetlin, and Winkler (2004) come closest to describing the risk-taking effect of different payoff schemes, as they separately show that the strategic choice of variability depends on the proportion of winners. When the proportion of winners is less than one-half, high variance strategies are preferred, whereas low variance strategies are best when the proportion of winners is greater than one-half. However, these papers stop short of describing the more complex rank-order payoff schemes found in real-world tournaments, most of which can be characterized by varying degrees of convexity. This paper picks up where they leave off to examine just how these more complicated payoff structures affect the equilibrium level of risk-taking.

Work by Audas, Barmby, and Treble (2004) looks at the effect of variation in corporate payoff structures. They use data from corporate hierarchies and show that as the difference in

pay from one level of promotion to the next increases, workers increase their effort. One of the underlying assumptions of the model presented here, as well as the model that motivated the research of Audas et al. (Lazear and Rosen (1981)), is that observable performance depends stochastically on effort. If that is true, then greater effort leads to a possibility of greater success and greater failure. Hence, the results of Audas et al. say something about the strategic choice of variability in corporate promotion tournaments, that workers increase the variance of outcomes as the spread in compensation from one level to the next increases. These results compare well with what we find.

In both content and methodology, this paper most closely resembles the work done by Gaba et al. (2004). Their study employs computational methods to find symmetric Nash equilibria of games in which agents choose the variance of their observable performance in a tournament. The game resembles a generic multi-round tournament wherein the players that finish with the highest performance are awarded a prize. Their model assumes that increasing one's variance is costless, and as mentioned, they find that agents should maximize variability when in a weak position (i.e. when a player is out-of-the-money) and minimize variability when in a strong position.

The differences between the current study and Gaba et al. are also significant, and it is important to mention a few here. First, rather than a generic game, this study offers a model of corporate promotion tournaments. Hence, the results on tournament design and risk-taking are applicable for managers as well as participants. These results give recommendations for the optimal design of corporate hierarchies in terms of the number of promotion levels and the salary structure. In a sense, these prescriptive implications for managers bring the literature on strategic choice of variability back to its roots in mechanism design [see Lazear and Rosen (1981) and Green and Stokey (1983)].

The second difference is that increased variability comes at a cost. In the context of the promotion contest, an agent may choose a high or low effort level, which determines stochastically her observable performance. Therefore, when putting forth more effort, there is a chance that the agent will be more successful, but there is also a chance that the agent will have nothing to show for it. This assumption draws a connection between the literature on strategic choice of variability and the literature on contests [see Nalebuff (1986) and Dixit (1987)]. One implication of this assumption is that the strategies do not imply a mean-preserving spread for the expected payoffs. Instead, more risk leads to a higher probability of promotion. Hence a comparison of expected payoffs depends on the salary structure, the number of rounds and the dynamic state of the game.

The third important difference is that while both papers use computational methods, this paper employs an evolutionary game theory approach rather than the iterative best response approach used by Gaba et al.. Using an evolutionary approach uncovers asymmetric as well as symmetric equilibria and offers an intuitive justification in terms of learning behavior. The drawback is that there are an infinite number of strategies, depending on tournament history and position, that can be used to adapt variance in a corporate promotion tournament. However, there is no way to account for all of these in an evolutionary setting. This point can be viewed as positive because the only option is to focus attention on a relatively small number of likely or intuitive strategies that have a strong foundation in the empirical literature as well as the lexicon of "folk knowledge."

2 The Model

2.1 Setup and Stage Game

The setting is modelled after a corporate promotion tournament. In each of T rounds, each of the N agents exerts some effort, and this effort corresponds stochastically to observable performance. Instead of choosing an effort level e_t , the agents choose the distribution from which to draw e_t . In other words, agents are playing mixed strategies $\sigma(e_t)$. To make things simple, there are only two choices of mixed strategies, σ_h and σ_l . They are:

$$\sigma_l = \begin{cases} 0 & \text{if } e_t < 0 \\ 1 & \text{if } 0 \leq e_t \leq 1 \\ 0 & \text{if } e_t > 1 \end{cases} \quad \text{and} \quad \sigma_h = \begin{cases} 0 & \text{if } e_t < 1 \\ 1 & \text{if } 1 \leq e_t \leq 2 \\ 0 & \text{if } e_t > 2 \end{cases}$$

Hence, σ_l is a uniform distribution over the interval $[0, 1]$, and σ_h is a uniform distribution over the interval $[1, 2]$.

With an exogenous probability, q , an agent is successful, and the agent's observable performance is e_t^* (realized effort). With probability $1 - q$, an agent is unsuccessful, and the agent's observable performance is 0. The probability q is interpreted as the difficulty of the task. It would be easy to allow for each level of the hierarchy to have different probabilities q_t . However, there is no obvious reason or method for assigning the different q_t 's, and it would also complicate the interpretation of the results. Therefore, $q_t = q$ for all t . With this setup, an agent's expected observable performance is $\frac{q}{2}$ when choosing σ_l and $1 + \frac{q}{2}$ when choosing σ_h .

At the end of each round, t , management makes its promotion decisions. To keep things simple, assume that after $t \leq T$ rounds, the N players are arranged in a pyramid according to cumulative performance up to round t . In other words, the $\frac{N}{2^{(t+1)+1}}$ agents with the highest cumulative performance are promoted to level $t + 1$ of the corporate hierarchy. The $\frac{2N}{2^{(t+1)+1}}$ agents with the next highest cumulative performance are posted (either promoted or remain) at level t of the corporate hierarchy. The $\frac{4N}{2^{(t+1)+1}}$ agents with the next highest cumulative performance are posted (either promoted, demoted¹ or remain) at level $t - 1$ of the corporate hierarchy, and so on until all $t + 1$ levels of the corporate hierarchy are filled. This scheme forms a pyramid of the following form:

Associated with each level τ of the hierarchy is a salary $\Pi(\tau)$. This is what an agent earns per period that she is in level τ of the hierarchy. The most basic assumption on Π is that $\Pi(\tau) \geq \Pi(\varsigma)$ whenever $\tau > \varsigma$ and $\Pi(\tau) > \Pi(\varsigma)$ for at least one pair of levels τ and ς .

In addition to the salary given by Π , an agent also expends some effort in each round, e_t^* . The agent bears the cost of this effort even if the effort is unsuccessful. Hence, an agent posted in level τ after round t earns $\Pi(\tau) - e_t^*$.

The distributions, σ_h and σ_l , correspond to high variance and low variance strategies respectively. When an agent chooses σ_h this clearly increases the probability of getting high salaries due to the increase in expected observable performance. Because the probability of higher

¹This can also be interpreted as a firing and rehiring of the agent at a similar firm at a lower level, due to the exchange of information between managers about the agent's cumulative performance.




Figure 1: The corporate hierarchy after t rounds.

observable performance increases by choosing σ_h in round t , the probability of higher salaries in all rounds after t also increases. This increase in the probability of higher salaries does not, however, increase the agent's probability of getting low salaries. Because every unsuccessful agent in round t has zero observable performance regardless of his or her effort, choosing a high variance strategy does not mean that unsuccessful agents that chose σ_h have higher probabilities of low salaries (in t or the future) than other unsuccessful agents.

Yet it is too simple to say that σ_h is an improvement over σ_l because the benefits of σ_h come with risk. In particular, the unsuccessful agent that chose σ_h has a lower expected payoff than an agent of equal cumulative performance that chose σ_l . The difference between σ_h and σ_l should then be interpreted as the difference between big risk and big rewards and small risk and small rewards.

One convenient interpretation of having strategy mixtures over effort is the following. Consider a sales contest in which each employee has the opportunity to pursue only one sale, and the promotion decision depends on relative sale sizes. The larger the sale, the more effort is required to make the sale successfully. However, when an agent chooses a sale to pursue, she does not know the exact size, only whether it is big or small. Therefore, choosing between σ_h and σ_l is like choosing between big and small sales, the exact size of which is *ex ante* unknown.

To model this simplified view of the corporate hierarchy analytically, it would require solving the following unsightly stochastic control problem:

$$\max_{\sigma(i)} \sum_{t=1}^T \left[\sum_{\tau=1}^t p(\Pi(\tau) | \sigma_t, \Omega_t, q) \Pi(\tau) - E(e_t | \sigma_t(i)) \right], \quad (1)$$

where

$$p(\Pi(\tau) | \sigma_t, \Omega_t, q) = p \left(\omega_{t+1}(i) \leq \Omega_{t+1} \left(\frac{2^{t-\tau} N}{2^{(t+1)} + 1} \right) | \sigma_t, \Omega_t, q \right),$$

$$\omega_{t+1} = f(\omega_t, \sigma_t),$$

$\sigma(i)$ is i 's T -round strategy, $p(\Pi(\tau))$ is the probability of getting salary $\Pi(\tau)$, σ_t is the collection of all agents' strategies in round t , ω_t is a vector for which the i th entry is agent i 's cumulative

Strategy Type	Description
0	always play σ_l
1	play σ_l in round 1 and when ahead thereafter
2	play σ_h in round 1 and σ_l when ahead thereafter
\vdots	\vdots
T	play σ_h until round $T - 1$ and σ_l when ahead. . .
T+1	always play σ_h

Table 1: Adaptive strategy types

performance up to and including round t , and Ω_t is a vector in which cumulative performances through round t are arranged in descending order. Also note that ω_t and Ω_t are random variables, the distributions of which are conditional on ω_{t-1} (itself a random variable) and σ_t .

By now it should be convincing that there is no hope of obtaining a closed form solution to 1. Therefore, to more fully understand the strategic choice of variance in this promotion tournament, it is informative to turn to an evolutionary setting.

2.2 Evolution

As discussed in the introduction, the folk wisdom that players in a tournament should take risks when behind and reduce risk when ahead is well documented in the empirical literature. While the intuition behind this type of strategy is clear enough, the mechanics of risk-adjustment are not at all understood. This is primarily due to the intractability of analytical models as demonstrated above. Here we turn to the power of an evolutionary model to glean results.

In a T -period tournament there is an infinity of possible histories. Therefore, even with only two actions, σ_h and σ_l , available in each round, agents have an infinite number of strategies to choose from for the T -period tournament. In order to be consistent with the empirical literature and common sense, we will only consider $T + 2$ strategies for a T -period tournament. These are presented in table 1.

Here “ahead” means being positioned in level $t + 1$ of the corporate hierarchy after t rounds of promotion decisions. For example, in the two round tournament; a type 0 strategist would play σ_l for both rounds; a type 1 strategist would play σ_l in the first round and σ_l in the second round only if ahead after the first round; a type 2 strategist would play σ_h in the first round and σ_l only if ahead after the first round; and a type 3 strategist would play σ_h in both rounds.

Alternatively, it is possible to allow agents’ strategies to depend directly on relative cumulative performance rather than current position. However, an agent’s current position is a direct result of her relative cumulative performance. Further, each agent’s current position in the corporate hierarchy is observable so that we do not need any additional assumptions regarding an agent’s knowledge of the other agents’ cumulative performance.

The strategy types are interpreted to be learned behavior rather than genetically programmed. Therefore, we can further interpret the replicator dynamic to be a learning dynamic rather than genetic evolution. Adopting this perspective avoids the awkward assumption that humans reproduce in proportion to their success in corporate promotion tournaments. The dis-

crete replicator dynamic is given below [see Weibull (1995) and Fudenberg and Levine (1998)]:

$$\phi_k = \frac{\phi_k U_k}{\sum_{j=0}^T \phi_j U_j} \quad \text{for all } k \in \{0, 1, \dots, T + 1\} \quad (2)$$

where ϕ_k stands for the proportion of type k strategists in the population and U_k is strategy k 's expected payoff.

Due to the difficulty of working with these types of tournaments in analytic form, evolution is simulated. Simulations provide insight into the comparative statics of the tournament setting, including the number of rounds, the payoff scheme and the task difficulty (q). Utilizing a simulation approach also focuses the analysis on a small number of simple and intuitive strategies, leaving aside the vast infinitude of all possible tournament strategies.

3 Results

The advantage of simulations is that they allow research to move forward in a particular area when analytical approaches are intractable, lack intuitive appeal or both. In such cases, simulations provide a way of mapping out general strategic considerations and also act as a “first cut” in providing both normative and predictive insights. While the results do not constitute an immediately generalizable theorem or proof, every effort is made to demonstrate the robustness of the simulation results to changes in parameters. With this in mind, the simulations identify many strong relationships between the tournament setting and the level of risk-taking.

Some of these results, while the result of straightforward intuition, are shown to be exceedingly general. For instance, the results support, not surprisingly, the notion that, in equilibrium, agents take on more risk when the probability of success increases. The intuition is, again, straightforward in that it is to be expected, holding all else equal that agents will put forth more effort when there is a greater chance of being rewarded. In fact, a convincing argument may be made that putting forth more effort cannot be characterized as taking on greater risk if it is in response to an increase in the probability of success.

Yet there are other very counterintuitive results which are also shown to hold quite generally. Not the least of which is related to the above intuitive result concerning an increase in the probability of success. The counterintuitive aspect comes in observing the following. As long as the payoff structure is sufficiently convex, low variance strategists (type 0) readily invade a population of high variance strategists (type $T + 1$), even when the probability of success is exceedingly high ($q = .95$). This is even more surprising when one considers that high variance strategies are monomorphically stable when $q = .75$ and the payoff structure is *less* convex.

Hence, following the logic that greater reward induces greater risk-taking is not sufficient when it comes to the convexity of the payoff scheme. In particular, the more convex the payoff scheme, the lower the level of risk-taking in equilibrium. The result is counterintuitive because intuition suggests that as the payoffs from higher levels of promotion increase relative to lower levels of promotion, there is an incentive to either do very well or do very poorly. There is not much difference between attaining level one and level two salaries, but there is a large difference between level $T - 1$ and level T salaries. Hence, it seems intuitive that such a structure would

encourage risk-taking in the same way that a convex utility function would. However, as the results show, this is not the case.

As mentioned, the simulations uncover several interaction effects and nonlinear relationships among the tournament parameters. For instance, when the payoff structure is convex, the level of risk-taking in equilibrium initially increases as the number of rounds increases, but then it reaches a maximum. After reaching the maximum level of risk-taking, the relationship between the number of rounds and the level of risk-taking becomes negative. Hence when plotted over tournaments of different lengths, the level of risk-taking exhibits a peak.

As the payoff structure becomes less convex, the relationship between tournament length and risk-taking begins to flatten out. Then when the payoff structure reaches linearity, so that the difference in pay between two consecutive levels of the hierarchy is constant, the relationship between tournament length and risk-taking becomes monotonic upwards. Interestingly, as the payoff structure becomes concave this relationship becomes nonlinear once again, exhibiting multiple peaks and valleys.

In terms of general strategic considerations, the model and results highlight the relative values of high, low and adaptive variance strategies under various conditions. It is interesting to note that adaptive variance strategists are not a part of an ESS for a wide range of parameter values. Instead of often finding positive proportions of adaptive variance strategists, the results reveal a different picture. Many tournaments converge to a state in which both high variance and low variance strategists maintain positive representation, yet adaptive variance strategists have died out.

When an adaptive variance strategy emerges as part of an equilibrium, very rarely is it an early adapter (a strategy that moves from σ_h to σ_l after being ahead early in the tournament). In fact, of the thousands of simulations and parameterizations examined not one produced stable non-zero proportions of a type that adapts before the halfway point of the tournament. Hence in a normative sense, the results do not support the strategy of adopting a risky strategy to establish an early lead and switching to a less risky strategy before the halfway point. As a general rule, if a player is ever going to play a low variance strategy, she should do so from the very start of the tournament or wait until late in the tournament to switch from high to low variance.

In what follows, features of the tournament that affect risk-taking are outlined. These features include the effects of the tournament length, the probability of success, the scale and convexity of payoffs, and the correlation of past success with future success. However, we first develop a measure of risk-taking in the population.

3.1 Level of Risk-Taking

Because this study is primarily interested in the incentives for risk-taking provided in the tournament setting, it is important to have a systematic method for measuring risk in a population. Therefore, we score the prevalence of risk takers in a tournament according to the distribution of strategy types in the following way:

$$R(T) = \sum_{t=1}^{T+1} \phi_t(t). \quad (3)$$

Call $R(T)$ the *risk score*. The risk score is the population average duration before switching from σ_h to σ_l . When $\phi_{T+1} = 1$, the risk score is $T + 1$, and when $\phi_0 = 1$, the risk score is 0. A risk score of $\frac{T+1}{2}$ implies that the average agent will switch to a low variance strategy if ahead halfway through the tournament. Using the risk score we can say that the level of risk-taking in the population increases as the distribution of strategists moves in the direction of strategists that wait a relatively long time before adapting from high variance (σ_h) to low variance (σ_l) or strategists that never adapt (type $T + 1$). Conversely, the risk score decreases as the population distribution moves in the direction of strategists that adapt early to play low variance or strategists that always play low variance (type 0).

3.2 Strategy Interaction

The first major question to be addressed is whether there are strategic interactions or not. If not, then the analysis of risk-taking in tournaments would boil down to a decision theoretic problem rather than a game theoretic problem. In that case, one needs only to find the optimal risk-taking strategy without regard to the strategies of other players and be done with the exercise. However, if the optimal strategy depends on what others are doing, then the problem is not so simple. Simulations reveal that the latter is in fact the case. Due to the minimal number of strategy types, the two-period version of the model provides the simplest approach to understanding the interactions among strategy types.

In a broad range of settings, strategy types exhibit strong non-monotonic evolutionary paths. Starting with a uniformly distributed population, these paths are often characterized by a rapid initial increase in the proportion of types other than type 0 (always play low variance). Hence the proportions of type 1, type 2, and type 3 increase at the expense of the low variance strategist. When ϕ_0 is low enough, ϕ_1 begins to decrease. Next ϕ_3 does the same. Over time, as ϕ_1 and ϕ_3 continue to decrease, the descent of ϕ_0 slows, and eventually its trajectory reaches a minimum very close to zero (sometimes as few as one or two individuals remain out of a population of 15,000). At the same time, ϕ_2 approaches (but does not reach) one. After ϕ_1 and ϕ_3 reach zero, ϕ_0 changes course and increases rapidly and ultimately overruns the population.

There are many tournament settings, like that illustrated in figure 2 where low variance strategists come close to dying off in well-mixed populations. However, with low proportions of type 1 and high variance strategists, these low variance strategists will invade a population of type 2. This result is interesting because, as stated above, early adapters like type 1 in a two-period tournament are never a part of an equilibrium that is reached from uniform initial conditions. However, type 1 does pose a strategic obstacle for low variance strategists as well as a strategic advantage for late adapters (like type 2).

The lesson is clearly that optimality depends on the population mix. For well-mixed populations, such as a uniform distribution, optimal play involves reducing variance if ahead after the second round. This is seen in the first 50 generations of figure 2. However, this strategy




Figure 2: Strategic interaction among types

loses to a low variance strategy if the two are competed head-to-head, as seen in the later generations of figure 2. The low variance strategy, in turn, loses to type 1 when they compete head-to-head, as seen in figure 3. Finally, figure 4 shows that type 1 loses to type 2.

Note that while it would be impossible to describe interaction effects for all population mixes and the entire parameter space, the above descriptions provide important insight about the limited applicability of popular strategy heuristics such as those captured by adaptive variance strategies. In particular, the optimality of playing a risky strategy until establishing a lead, then switching to a less risky strategy depends highly on what the other players are doing. In the example above, using the adaptive variance approach is optimal in a well-mixed population and certain pairwise competitions, but suboptimal when paired against low variance strategies.

Next, we vary the tournament length to examine the effect on risk-taking.

3.3 Tournament Length

The length of the promotion tournament has implications for risk-taking behavior and optimal adaptation. Yet these implications differ when comparing different types of payoff schemes. When comparing risk-taking across different tournament sizes, it matters greatly whether we control for the change in salary from one level to the next or the difference in salary between the lowest and highest levels.

To be systematic in comparing the amount of risk-taking in a tournament of size T it is convenient to use the risk score. However, the problem with the risk score is that, while it




Figure 3: Low variance vs. type 2

allows us to compare the prevalence of risk-takers in tournaments of equal length, it does not allow for comparing risk across tournaments of different lengths. Hence, the *length-adjusted risk score* normalizes the risk score by tournament length.

$$\mathcal{R}(T) = \sum_{t=1}^{T+1} \phi_t \frac{t}{T+1}.$$

An interpretation of $\mathcal{R}(T)$ is the population average duration before switching from σ_h to σ_l , normalized for tournament length. Hence, like the risk-score, the length-adjusted risk score gives us a measure of risk-taking in a given population. In what follows, the length-adjusted risk score is employed to quantify the difference in risk-taking across tournaments of different lengths.

As in all games, the payoffs are essential. Hence, we take care to categorize our comparisons of different tournament lengths according to the payoff structure. The first and simplest category is when the difference between salaries at two consecutive levels are controlled across varying tournament lengths. Here payoffs are awarded so that each level of the hierarchy gets a salary exactly δ units greater than the previous level. With this setup, regardless of the value of δ , the equilibrium level of risk-taking increases with tournament length. This implication holds for all $\delta > 0$, values q (probability of success), and a wide range of initial conditions.

The following figure reports length-adjusted risk scores for tournament lengths from two to eight rounds after the population has converged to an evolutionarily stable state from uniform initial conditions. The payoff structure employed is such that $\delta = 2$. Hence, each level of the hierarchy, τ , gets 2τ , and the top level in an eight-round tournament gets 16, while the top




Figure 4: Type 1 vs. type 2

level in a two-round tournament gets just 4.

However, as suggested, increasing the length of the tournament comes with an increase in the difference $\pi(T) - \pi(1)$. So it's difficult to say whether the apparent positive relationship between risk-taking and tournament length is true without controlling for this. To do so, first fix $\pi(1) = 2$ and $\pi(T) = 2 + \Delta$. Next, set:

$$\pi(\tau) = \pi(1) + (\tau - 1) \frac{\Delta}{T - 1}. \quad (4)$$

This form standardizes the difference $\pi(i) - \pi(j)$ for all consecutive i and j such that $i > j$. Figure 6 compares the baseline from figure 5 with the standardized payoffs. Using this approach highlights the effect of controlling $\pi(T) - \pi(1)$ on the relationship between tournament length and risk-taking. In the baseline case from figure 5, the level of risk-taking increases with tournament length (until it reaches one). However, when controlling for $\pi(T) - \pi(1)$, the level of risk-taking initially increases along with tournament length, but then it levels off. Another difference is that in the baseline case, the initial increases in risk-taking are more gradual than in the controlled case. Hence, while the relationship between tournament length and risk-taking is essentially a positive one for payoffs that increase by a fixed amount, controlling for the difference $\pi(T) - \pi(1)$ dilutes the strength of that relationship.

Next we vary the level of convexity and concavity of the payoff scheme to examine the effect on risk-taking.




Figure 5: Length-adjusted risk score and tournament length

3.4 Payoff Convexity

As is the case with all games, the optimal strategies in promotion tournaments are highly sensitive to the payoff structure. Of particular interest when looking at risk-taking in tournaments is the convexity of the payoff structure. As already touched on, the effect of increasing the convexity of the payoff structure is counterintuitive. In general, the more convex the payoff structure, the lower the level of risk-taking in equilibrium.

Alternatively, it is possible for the difference between the salaries of one level and the next to be systematic, but not constant. For example, the change in salary from one level to the next can increase by some constant factor (similarly, the salary itself can increase by a constant factor). The results obtained when the change in salary from one level to the next doubles resemble the results from the simpler case in the previous section where the difference is constant. In the table below, results are reported from simulations in which the difference in salary from one level to the next increases by a factor of two. We fix $\pi(1) = 2$ and then allow $\pi(2) - \pi(1) = 1$ so that $\pi(3) - \pi(2) = 2$ and so on.

Similar to the results obtained when the increase in payoffs is constant, these convex payoffs produce a positive monotonic relationship between tournament length and risk-taking. However, there are two differences in the payoff schemes, the convexity of the payoffs and the difference $\pi(T) - \pi(1)$. Therefore it is not obvious which of these changes is driving the observed relationship between tournament length and length-adjusted risk score. Is it the convexity of the payoff scheme or the difference $\pi(T) - \pi(1)$?

In order to compare the effect of convexity on risk-taking, we must control for the difference $\pi(T) - \pi(1)$ in a way that allows us to vary convexity systematically across tournament lengths.

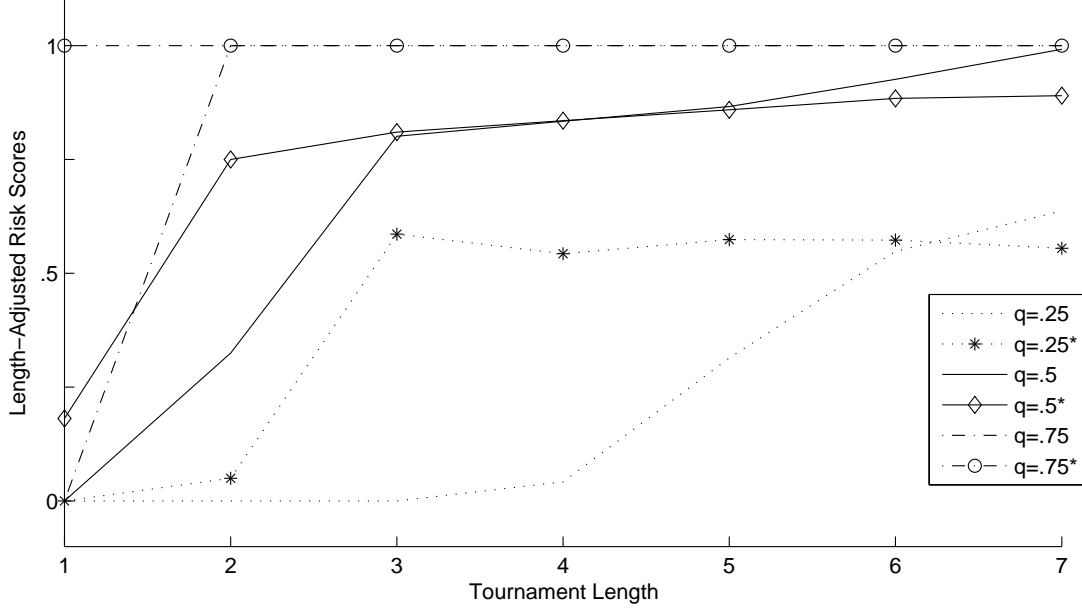


Figure 6: Baseline payoffs (q) and payoffs when $\pi(T) - \pi(1)$ is controlled (q^*).

Doing so gives us the payoff structure from equation 4. This form can easily be adapted to the following more general form:

$$\pi(\tau) = \pi(1) + \sum_{t=2}^{\tau} \alpha^{t-2} \delta(\alpha). \quad (5)$$

Here $\pi(1)$ is fixed and $\tau > 1$. Note that $1 > \alpha > 0$ produces a concave payoff structure and $\alpha > 1$ produces a convex payoff structure, whereas $\alpha = 1$ reduces to equation 4. The constant $\delta(\alpha)$ is simply the difference $\pi(2) - \pi(1)$. Once α is chosen, $\delta(\alpha)$ is computed as:

$$\delta(\alpha) = \frac{\Delta}{\sum_{t=0}^{T-2} \alpha^t}.$$

With this more sophisticated setup, it becomes easy to examine the effect of various payoff convexities on the level of risk-taking. Setting the difference $\pi(T) - \pi(1) = 10$, we produce the payoff schemes in figure 7 below. Note that $\alpha < 1$ produces concave payoff schemes while $\alpha = 1$ produces payoff schemes with constant increase in salary and $\alpha > 1$ produces convex payoff schemes.

Comparing the results from the normalized linear payoff schemes ($\alpha = 1$) against results from simulations of normalized convex payoff schemes ($\alpha = 2$) illustrates the effect of convexity on risk-taking. The tournament has four levels of promotion, and the probability of success (q) is varied. It is evident that a linear payoff scale ($\alpha = 1$) maximizes risk-taking (effort) for a broad range of values q . Increasing convexity or concavity leads to a reduction in the equilibrium level of risk-taking in the four level case. However, there is no reason to expect this exact relationship to hold across tournaments of different lengths. In general, for all convex payoff structures the

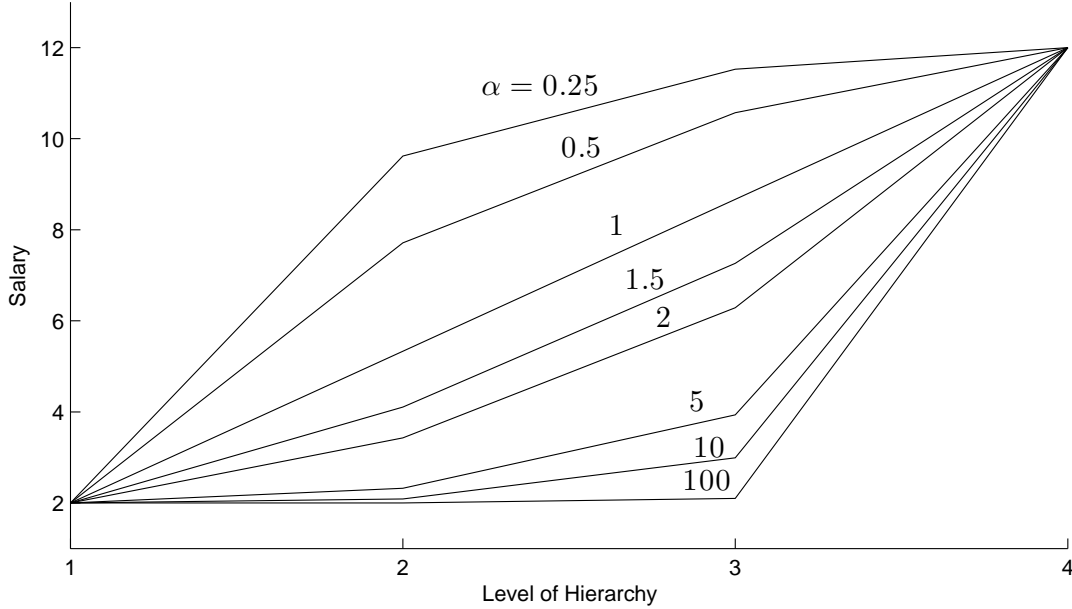


Figure 7: Convexity of payoffs in a four-level tournament

degree of risk-taking initially increases as the tournament length increases, but it then reaches a maximum, after which there is a negative relationship between risk-taking and tournament length. However, as already suggested at the start of this section, the convexity of the payoff scheme does have implications for the relationship between tournament length and risk-taking. In particular, as the convexity of the payoff scheme is increased, the length of tournament that maximizes risk-taking decreases.

In terms of tournament design, this means that a manager can manipulate the equilibrium level of risk-taking in a promotion tournament by adjusting the number of promotion levels. For instance, suppose a manager wants the entry level position to be paid $\underline{\pi}$ and the top level position to be paid $\bar{\pi}$, then this manager can optimize the amount of risk-taking by choosing the number of intermediate promotions and the convexity of the payoffs appropriately. Figure 9 displays this relationship for various choices of α . From this we see that $\alpha = 1$ maximizes the level of risk-taking for two, four, six, seven and eight round tournaments, whereas more convex payoff schemes maximize risk-taking for three round tournaments and $\alpha = 1.5$ achieves the maximum risk-taking for five round tournaments.

4 Conclusion

Although the prevalence of adaptive variance strategies are well-documented in the economic literature, many aspects of the strategic choice of variability are not. In particular, relatively little is known about the effect of tournament design on the level of risk-taking in equilibrium. In the context of corporate promotion tournaments, the tournament design includes the num-




Figure 8: Convexity of payoffs and length-adjusted risk scores

ber of promotion levels and the associated salaries. This study looks at the impact of these features of tournament design on the strategic choice of variability and uncovers several strong relationships.

The first lesson is that early adaptation is very rarely a component of optimal play. Only when matched against certain strategies, such as overly conservative low variance strategies, does an early adaptation emerge in equilibrium. But for well-mixed (or uniformly distributed) populations, early adapters lose to more aggressive strategists or less aggressive strategists or a mixture of both. In fact, the vast majority of equilibria are comprised of some combination of low variance strategists, very late adapters, and high variance strategists. Hence, if a player elects to adapt the variance of her strategy, she should wait until the last minute to do so. With that said, the finding that dominant strategies are rare warrants another mention. Strategic interactions are strong, and so optimal strategies depend highly on the population mix.

Another important lesson is from the perspective of tournament design. Managers should be aware that the level of risk-taking can be optimized with respect to the number of levels and the convexity of the salary structure. For many situations, especially as the number of promotion levels increases, the salary structure wherein each level of the hierarchy earns $\$ \delta$ more than the level lower maximizes risk-taking and effort. This approach to studying the strategic choice of variability reintroduces the important insights of mechanism design. There is ample room for contributions to the understanding of optimal design with respect to risk-taking (effort) in many tournament settings, especially in the realm of elimination tournaments.

Due to the highly complex nature of this topic in game theory research, computational and evolutionary tools continue to be of essential use. Additionally, results vary greatly depending on the specific situation modelled and are not easily generalized to a broad class of games. This means that each tournament setting deserves individualized research attention.

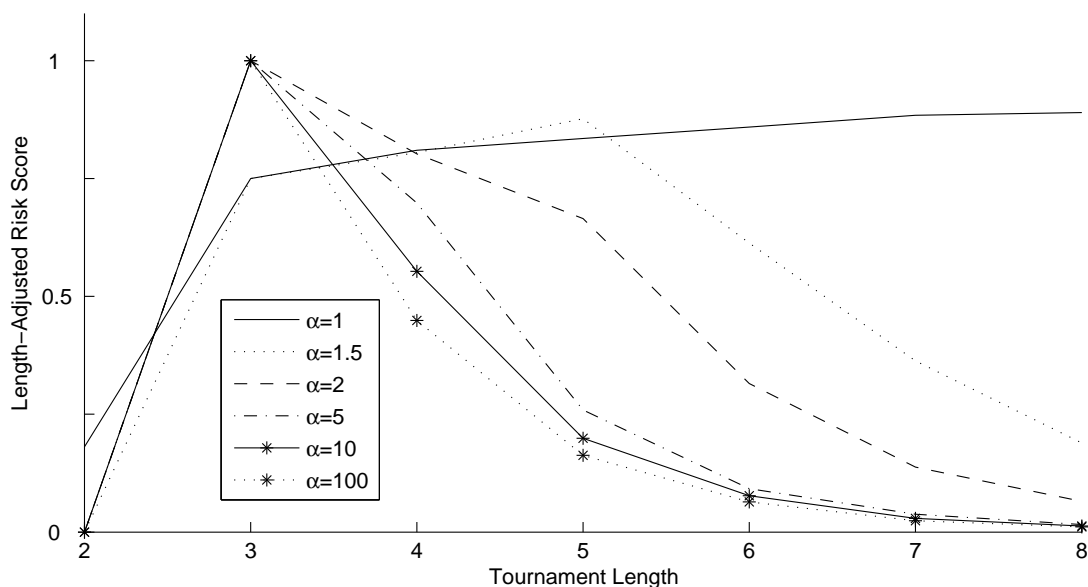


Figure 9: Length-adjusted risk score and tournament length with convexities

References

- R. Audas, T. Barmby, and J. Treble. Luck, effort and reward in an organizational hierarchy. *Journal of Labor Economics*, 22(2):379–395, 2004.
- M. Bognanno. Corporate tournaments. *Journal of Labor Economics*, 19(2):290–315, 2001.
- S. Bronars and G. Oettinger. Effort, risk taking and participation in tournaments: Evidence from professional golf. Working Paper, University of Texas, 2001.
- K. Brown, W. Harlow, and L. Starks. Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *Journal of Finance*, 51:82–110, 1996.
- L. Cabral. R&d competition when firms choose variance. *Journal of Economics and Management Strategy*, 12:139–150, 2003.
- J. Chevalier and G. Ellison. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy*, 105:1167–1200, 1997.
- A. Dixit. Strategic behavior in contests. *American Economic Review*, 77:891–898, December 1987.
- D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- A. Gaba, I. Tsetlin, and R. Winkler. Strategic choice of variability in multiround contests with handicaps. *Journal of Risk and Uncertainty*, 29(4):143–158, 2004.

- A. Goriaev, F. Palomino, and A. Prat. Mutual fund tournament: Risk taking incentives induced by ranking objectives. *CEPR, Tilburg University, Working Paper*, 2001.
- J. Green and N. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–364, June 1983.
- C. Grund and O. Gürtler. An empirical study on risk-taking in tournaments. *Applied Economics Letters*, 12(8):457–461, June 2005.
- E. Lazear and S. Rosen. Rank order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864, 1981.
- B. Nalebuff. Prizes and incentives: Towards a general theory of compensation and competition. *Bell Journal of Economics*, 13:21–43, 1986.
- J. Weibull. *Evolutionary Game Theory*. MIT Press, 1995.
- Y. Yumoto. *Risk Taking in Promotion Tournaments*. PhD thesis, North Carolina State University, November 2003.