Prize Structures, Gender, and Performance in Competitive Marathon Running

Dan O'Toole
Economics 396b / Research Seminar
Advisor Prof. Jeff Frank
April 30, 2009

Table of Contents

I. Introduction……………………………………………………………………2

II. Tournament Theory…………………………………………………………3

III. Tournaments and Sports Economics……………………………………..5

IV. Empirical Investigations of Tournament Theory………………………...7

V. Empirical Investigations of Road Racing……………………………….8

VI. Gender and Competitive Environments………………………………...11

VII. Data and Methodology…………………………………………………13

VIII. Regression Models and Results………………………………………...16

IX. Commentary……………………………………………………………20

X. Summary and Conclusions……………………………………………...22

XI. Appendix I………………………………………………………………24

XII. References………………………………………………………………26

I. Introduction

Well before the advent of the World Marathon Majors Series in 2006,[1] competitive marathon running has been a thriving business, with race organizers looking to design events that exhibit elite runners at their most competitive. Designing a marathon forces race organizers to deal, however, with the idiosyncrasies of the distance; the marathon is a far different animal from the local road races that can be found throughout the country on a given Saturday morning. With the physical toll that racing a fast marathon (26.2 miles) takes on a runner's body, and the amount of training necessary to compete in a marathon at a high level, athletes typically only compete up to twice in a given year. One crucial result of this feature of marathon running is the assumption that elite marathoners (more so than slower runners, and runners competing at shorter, more frequently-raced distances) operate with income-maximization in mind. Given this assumption, the purpose of my paper is to test the responses of competitive marathoners to differences in prize level and structure using data from the 12 largest United States marathons in 2005. From a theoretical perspective, my analysis is enabled by the theory of tournaments, a concept from personnel economics that has frequently been applied empirically to athletic contests. Modeled first by Lazear and Rosen (1981), tournament theory considers labor contracts that are based not simply on a worker's absolute physical output, but rather on their rank relative to other workers in an organization. Running, as well as a number of other individual sports, thus provides a fertile testing ground for tournament theory in action: competitors are awarded based primarily on their final ranking in the contest. As the prize structure is known in advance and fixed throughout the competition, and because it is easy to operationalize performance (e.g. final time, number of

[1] The World Marathon Majors is a series of five international marathons with a combined purse of $1 million for males and females. It was launched in January 2006 with the intent of increasing the fitness incentives of elite runners and, consequently, raising international awareness of marathon running.

strokes in golf, etc.), it becomes possible to test empirically the effects of different prize structures on contestant performance.

But unlike similar empirical work that has been done on road racing, I wish to incorporate another key variable: gender. With a sample comprising performance data for both males and females, I will attempt to compare the effects of prize structures on performance across the genders. Much empirical work has been conducted on gender differences in competitive environments, and I will be introducing the arena of the marathon to those investigations. The outline of my paper is as follows. After reviewing the key principles of tournament theory, I will discuss its extension to athletic contests. From this theoretical base, I will then review the existing literature that has applied tournament theory to sports, focusing more heavily on that which focuses on running, and also the existing literature on gender differences in competitive environments. The subsequent section presents my regression models and results, and the final two sections are a commentary on the results and summary and conclusions.

II. Tournament Theory

What fundamentally differentiates tournament theory from other measures of performance—whether in a corporate organization or a sporting contest—is the dependence on a relative rank instead of an absolute standard. Lazear and Rosen model tournament theory from a personnel-economic perspective, and it is largely from this seminal paper that the theory has been considered and applied. The economists distinguish their work from that which analyzes incentive payment schemes that are linked to output (piece rates), analyzing instead compensation schemes that are linked solely to the rank-order of an individual in an organization. Part of the rationale for this mode of compensation, Lazear and Rosen write, has to do with the information costs of measuring worker performance. In firms in which it is less costly to measure relative rank than to measure worker productivity, compensation schemes based on the former measure are favored. The economists also seek to examine the relative efficiency of these schemes. They demonstrate that, under the conditions of competitive labor and product markets, and a Nash equilibrium solution standard, rank-order compensation can lead to an equally efficient allocation of productive inputs as a compensation scheme based on absolute productivity. There exists, therefore, some set of prizes for relative performance (i.e. rank) that result in efficient labor effort and, by extension, efficient investment in human capital on the part of a firm.

From this framework, there are three basic principles of tournament theory. First, prizes are fixed in advance, and known in advance by the contestants. Second, larger compensation spreads between positions in a hierarchy induce those (initially) at the lower levels to put forth more effort. And third, there is an optimal compensation spread. While, as the second principle marks, larger spreads induce more worker effort, at a certain point costs exceed benefits; this marginal cost of additional worker effort is only worth incurring from the standpoint of the firm if it is exceeded by the marginal revenue product realized from additional worker output.

In addition to these three central principles, Lazear and Rosen also attempt to incorporate several other variables that affect the behavior of contestants in tournaments. One of these is a variable measuring luck and randomness in a tournament setting; contestant performance is a function of both effort and this variable. If luck significantly affects a contestant's probability of winning, contestants will try less hard to win, as the effort they expend in the contest has a smaller effect on winning. In such risky environments, in which exogenous forces are likely to influence outcome, Lazear and Rosen write that there must be larger spreads between ranks (larger marginal increases in compensation) so that contestant effort does not wane in the face of randomness and uncertainty. Another important variable in the economists' model is the number of 'winning' slots present. Something like a 'winner-take-all' situation (e.g. if a firm has 60 vice-presidents all competing for one president position) would diminish contestant effort, as many contestants would put forth little effort, thinking the promotion to be unlikely. A seemingly opposite situation, in which 'winning' is effectively guaranteed (e.g. 60 vice-presidents and 60 president positions), would produce a similar effect, however. If winning were guaranteed in the tournament, contestants would have little incentive to exert substantial effort.

III. Tournaments and Sports Economics

Though Lazear and Rosen theorize tournament theory as related to personnel economics, the framework they provide maps on naturally to sporting contests. The work of Demsetz (1995) offers a useful preliminary framework for considering sports through this lens; extensive empirical work has also been done, which I will review in the two next sections. The economist writes that professional sports provide a viable application because compensation is awarded primarily based on the final rank-order of contestants. Demsetz does not contend, however, that it is because of the information costs of measuring and monitoring productivity that this compensation scheme is favored. He writes that, in many sporting contests, it would be equally feasible to award compensation based on absolute productivity (e.g. fewness of strokes in golf or finishing time in a race) as on rank-order, as the former is easily measured. Still, Demsetz acknowledges reasons why rank-order schemes are favored in professional sporting contests. The economist makes the point that productivity is sporting contests is ultimately tied to consumer satisfaction in their witnessing of the event, rather than tied narrowly to the physical output of contestants. With the maximization of consumer satisfaction in mind, the compensation schemes of sporting contests should be designed to induce maximal, or 'all-out,' effort by contestants.

Given a specific set of preconditions (level playing field, enforceable rules, and well-matched or relatively homogenous contestants), a rank-order compensation system will induce, Demsetz writes, the highest level of joint productivity (i.e. consumer satisfaction) by the contestants. If compensation is not tied to absolute physical output, then a rank-order compensation scheme will bring out high levels of effort from contestants because even a well-performing second place contestant will earn significantly less than the winner. In other words, small differences in physical output result in only small differences in compensation under a non-rank-order system, but will likely result in larger differences in compensation under a system considering only the relative rank of contestants. In a close race, for example, it is likely that two competitors will exert maximal effort only if the opportunity cost of finishing second is sufficiently high (i.e. if the difference in first prize and second prize is sufficiently large; cf. Lazear and Rosen). With non-rank-order compensation schemes, these opportunity costs (i.e. differences in compensation) tend to be lower. Demsetz's ultimate point, therefore, is as follows: if the highest level of joint productivity (again defined as consumer satisfaction) is the objective of a contest, then a rank-order compensation scheme should be put in place.

IV. Empirical Investigations of Tournament Theory

As sporting contests provide much observable data regarding performance and compensation, it is natural that there has been much empirical work applying tournament theory to athletics. Many of these investigations have taken the form of researchers' attempts to understand the effect of different incentive or compensation structures on the performance of individual athletes in competition with one another. As Szymanski (2003) notes, most individual sports fall into the model of a head organizer whose has the objective of eliciting some sort of investment or effort from competitors who may, as a result of that investment or effort, win a prize. Furthermore, it is the objective of the organizer to design an incentive scheme that maximizes the entrants' effort, as that effort is the prevailing measure of consumer satisfaction (and customer satisfaction can translate, in turn, an enhanced reputation for the event director, which can translate, in turn, to increased media attention, corporate sponsorship, revenue for the host site, etc.).

What researchers have largely tried to test in sports settings, then, are the impact of prize structures on the performance of competitors. In a widely-cited paper, Ehrenberg and Bognanno (1990) consider the incentive effects of tournaments using data from the 1984 Men's PGA tour and the 1987 European Men's PGA Tour. As both the incentive structure (prize distribution) and individual output (players' scores) are easily observable, these tournaments offer an effective, non-experimental arena in which to test the contest theory model. The researchers' chief finding is that scores tend to be lower (i.e. performance higher) when the total prize fund is larger. Furthermore, Ehrenberg and Bognanno considered the effect of marginal winnings (i.e. change in winnings from a one-place improvement in overall rank) on players' final round scores. The researchers' hypothesis follows from Lazear and

Rosen's principle of the positive relationship between worker effort and the compensation spread between successive ranks. Since in golf (and in other sports like marathon running), the prize spread is diminishing with respect to rank (e.g. the difference between first and second is much greater than between eighth and ninth; see Appendix I), Ehrenberg and Bognanno hypothesize that player effort will be higher (scores lower) in the final round the higher that player's placing at the beginning of the round. The data confirm this hypothesis, as players facing higher marginal returns to effort were observed to perform better in the final round of tournaments. Becker and Huselid (1992) also attempt to test tournament theory in the setting of sports, using data from the National Association for Stock Car Auto Racing (NASCAR) and the International Motor Sports Association (IMSA). Using panel data controlling for driver ability, the researchers find that variation in absolute spread in prize money between higher and lower finishers produces a significant positive effect on driver performance. As the spread increases, though, driver performance increases at a diminishing rate.

V. <u>Empirical Investigations of Road Racing</u>

Tournament theory has also been used in empirical considerations of distance running, the area of athletics which my own research will address. Lynch and Zax (2000) test the incentive effects of tournaments in a sample of runners competing in distances from 5 kilometers (3.1 miles) to the marathon (26.2 miles). Given their panel data from a large number of races (multiple observations for an individual over the period of a year), the researchers were able to an individual runner's performance in multiple races, with the aim of comparing performance over different prize schemes. Lynch and Zax conduct this comparison through two different models of regression. Without controlling for the runner's pre-race ranking relative to the other runners in a given race (a "point" ranking based on personal-best times over each distance raced), the researchers found that the greater the sum of prizes, the faster the winning performance. But when runner ability was controlled for through fixed-effects, the incentive effects were much smaller. Lynch and Zax posit that this finding might be the result of sorting. For race distances longer than 10 kilometers, the researchers' test for joint significance of runner dummy variables is highly significant. They hypothesize, therefore, that runner variables are important in explaining the variation in times; it is not just the characteristic of a particular race, then, that explains finishing times. According to this concept, runners would select races based, in part, on the level of prize money they hope to win. Hence, races with greater total prizes would attract faster runners, and the winning performance would reflect such selection tendencies.

In a paper more in line with my own project, Frick and Prinz (2007) consider, unlike Lynch and Zax, only the marathon distance, testing a variety of hypotheses related to tournament theory. The researchers frame their study, however, as critically different in its data set. Frick and Prinz write that aggregating data from races that range from 5 kilometers to the marathon produces results that are likely unable to be interpreted. First, race tactics change dramatically across distances; viewing those competing in a 3.1 mile race and a 26.2 mile race as constituting a homogenous population is misguided. Furthermore, because of the physical toll a 26.2 mile race takes on a runner's body, competitive marathon runners are likely to race sparingly, rarely more than twice in a given year. This feature of the marathon has at least two important consequences. First, because of the infrequency with which the distance can be raced, marathoners will enter marathons only if they are able to perform to their maximum capacity. If the races from which data is gathered have large enough fields (which is to say, are high-profile and competitive enough), then individual performances will be more homogenous than performances gathered from a range of race distances. As a result, it becomes easier to measure the effect of prize structures on runner performance: variables like natural ability levels and risk behavior are fairly consistent across observations (at least to a greater degree than when considering a range of distances). Second, the length of the physical recovery period from a marathon induces runners to allocate their physical resources (i.e. effort) more wisely—that is to say, in a manner more compatible with income maximization. For this reason, competitive marathoners must consider both their potential income from completing a race and the opportunity cost (foregone prize money from future races) of a poor—but still completed—race. Yet again, limiting data to performances from the marathon distance creates a more homogeneous sample; those who complete a marathon can be seen to behave with income-maximization in mind.

Frick and Prinz collect data on men's performances in 13 different major city marathons held in the United States, Britain, and Germany from 1983-2001, considering the total purse, additional bonuses for absolute (time) performances, prize distribution, and prize differences between adjacent ranks as parameters. The researchers find that each race's monetary awards generally had the predicted effect of improving overall performance. The higher the total purse and the more unequal its distribution, the faster the finishing times. The data show that doubling the average prize leads to a reduction in finishing times by 1.5%; that a larger prize spread leads to faster races (or higher overall effort, a finding very much in line with tournament theory); that smaller prize differences between adjacent ranks produce slower races; and that doubling bonus payments leads to a reduction in finishing times by 3%.

Trying to determine the closeness or 'excitement' level of a race, Frick and Prinz regress a variable they refer to as "temporal backlog": the time gap behind the winner of finishing positions 2-5. They find that a larger total purse produces smaller temporal backlogs, or a more tightly-clustered field. Finally, having collected data on marathons' bonuses for absolute performance, the researchers find that such a time-based incentive mechanism did in fact spur runners to maintain high effort levels.

VI. Gender and Competitive Environments

In analyzing data from only male performances, though, Frick and Prinz ignore a potentially productive line of investigation: gender differences in competitive environments. My own research will attempt to address this set of questions as they pertain to marathon running; before I do so, however, it is first important to review the existing literature on the subject. Gneezy and Niederle (2003) attempt to explain the gender differences in competitive high-ranking positions by conducting experiments in laboratory settings. Working against the literature that has attributed these differences to workplace discrimination (e.g. Black and Strahan 2001), the researchers conduct controlled experiments to test the hypothesis that men and women react differently to competitive incentive schemes. Using the activity of solving mazes, the researchers find that when participants were paid a noncompetitive piece rate (based only on their nominal performance), differences in performance between genders were statistically insignificant. When participants are paid competitively (i.e. based on their performance in solving mazes relative to other participants), however, the researchers observed a significant increase in performance relative to the benchmark in male participants, but not of female participants. The researchers also found that in mixed-gender environments, men significantly outperform women.

Gneezy and Rustichini (2002) also conduct experiments testing gender differences in competiveness, though among children aged nine to ten. The researchers attempt to test, in part, whether same-gender environments have the same effect on performance as mixed-gender environments. Having all the children run separately for a short distance, the researchers obtain a benchmark measure of performance for each participant. The children are then matched in pairs according to their initial performance (pairs were both mixed-gender and same-gender), and made run a second time. In this second round of races, the mixed-gender pairs show improved performances, though the effect of competition on males was significantly higher than on females. In the male-male pairs, performance also improves. In the female-female pairs, though, participants' performance did not improve significantly. The researchers' work suggests, therefore, that females are less motivated to perform in competitive environments than males, and that, for females, intra-gender competition also diminishes this motivation (taking into account, of course, the age of the participants and the physical rather than mental nature of the task, among other caveats).

In a similar vein, Niederle and Vesterlund (2007) test whether males and females differ with respect to their selection into a competitive environment. Somewhat similar to Gneezy and Niederle's (2003) work, the researchers give participants in a laboratory setting a physical task to perform, first under a noncompetitive piece rate and second under a competitive tournament incentive scheme. Though absolute performance does not differ between genders, men are much more likely to select the tournament scheme for their next task than women (73% of men compared to 35% of women). Controlling for absolute performance, overconfidence, and risk, the researchers determine that the differences of entry into tournaments can be attributed to a real difference in preferences between the genders.

VII. Data and Methodology

The data used in my own investigations come from the 12 largest United States marathons (4,000 total finishers or more) that awarded prize money in 2005; data are collected, of course, for both genders (total n = 216). For descriptive summary statistics for each marathon in the sample, see Tables 1a (males) and 1b (females). As these tables show, prize structures (both prize level and prize distribution) are identical across genders; Appendix 1 shows graphs of selected prize distribution schemes for the marathons of the sample. Seeking, first, to explain the variation in runner performance (operationalized as an individual's finishing time), the parameters to be considered are a runner's gender, the average prize money per prize-eligible runner, the average marginal prize spread between adjacent places (i.e. average of the marginal returns to finishing one place higher and one place lower), and finally, gender interaction terms for the two incentive variables. My other regression model in this study seeks to draw conclusions more marathon-level (competition) than individual-level (runner). Using the same set of parameters, the other performance-related variable I am seeking to explain is what I am calling "time gap": the average time behind the winner of runners finishing in places 2 through 5 in a given race. Somewhat more comprehensive than individuals' finishing times, the time gap variable can be seen as a measure of the closeness or excitement level of a given race. As this measure is likely tied to spectator/consumer satisfaction (which Demsetz would mark as the ultimate productivity of the contest),

Table 1a: Descriptive Statistics of the Sample (Males)

Marathon	Finishers	Winning Time	Time Gap[2]	Prize Money[3]	Prize Depth
New York City	25279	2:09:30	58	236	10
Chicago	18677	2:07:02	22	270	5
Los Angeles	12409	2:09:35	148	73.5	10
Boston	10894	2:11:45	95	242.5	15
Rock 'N Roll (CA)	7467	2:09:17	130	79.1	10
Las Vegas	4829	2:11:57	157	46.9	10
Twin Cities (MN)	4685	2:18:28	198	73.5	10
Rock 'N Roll (AZ)	3727	2:14:24	120	50	10
Grandma's (MN)	4344	2:13:18	75	36	10
Philadelphia	3824	2:21:02	186	7.75	5
Houston	3706	2:14:50	280	48	8
Austin	3024	2:12:38	169	29	5

Table 1b: Descriptive Statistics of the Sample (Females)

Marathon	Finishers	Winning Time	Time Gap	Prize Money	Prize Depth
New York City	12318	2:30:55	684	236	10
Chicago	14326	2:21:25	290	270	5
Los Angeles	7576	2:26:11	209	73.5	10
Boston	6655	2:25:13	304	242.5	15
Rock N' Roll (CA)	8467	2:30:55	60	79.1	10
Las Vegas	3357	2:31:54	21	46.9	10
Twin Cities (MN)	3068	2:40:21	157	73.5	10
Rock N' Roll (AZ)	3439	2:32:51	1071	50	10
Grandma's (MN)	2541	2:28:43	405	36	10
Philadelphia	2063	2:43:07	348	7.75	5
Houston	2021	2:32:27	468	48	8
Austin	1930	2:31:01	969	29	5

[2] Average time in seconds behind winner of the 2nd - 5th finishers in the race.

[3] In 1000s (2005 US$).

it is important to attempt to trace its potential relations to incentive structures. For detailed definitions of these key variables, see Table 2.

Table 2: Definitions of Key Variables

Variable	Definition
LNTIME	log of runner i's time in seconds in race j
LNTIMEGAP	log of average time in seconds behind finisher of runners finishing in places 2-5 in race j
LNSUMPRIZE	log of average prize money per prize-earning place per gender in race j
MARGDIFF	log of the average of runner i's returns ($) to finishing one place higher and one place lower in race j
GENDER	gender of runner i in race j

By marking out a sample that consists of only runners eligible for prize money, and only such runners in the largest US marathons, I have attempted to generate a relatively homogenous set of observations. Like Frick and Prinz, by limiting my sample to only the marathon distance, I have tried to ensure that my observations reflect elite runners who behave with income-maximization in mind. As noted before, the physical toll that racing a marathon brings, and the scarcity with which a runner can compete at the distance, tends to ensure that runners who enter and complete the race do so at maximal or close to maximal fitness levels. Given this benchmark of commitment to the marathon being raced, it therefore becomes easier to operate under the assumption of income-maximization. This assumption allows us to conceive each runner observed as constitutive of a homogenous group—similarly inclined with respect to physical ability and risk attitudes. Lacking, as I do, an extensive panel data set (with multiple observations for many individuals), this feature of the sample enables us to measure the effects of prize structure on performance without having to control for a range of individual ability levels. Finally, while marathons are all the same distance (26.2 miles), the courses on which they are run are not identical. Fortunately, the winning times of the marathons which make up my sample are close enough (within just over one standard deviation of one another) to rule out variation in course layout as correlated to variation in finishing times. Including a variable such as a marathon's course record (i.e. the faster time ever recorded on a given course) to control for these differences would seem logical; but because prize structures change from year to year, this variable and the incentive variables would show collinearity. For this reason, it was omitted from my eventual model.

VIII. Regression Models and Results

From the parameters listed in the previous section and defined in Table 2, my regression models are of the following two forms:

1) $LNTIME_{ij} = \beta_0 + \beta_1\, LNAVGPRIZE_j + \beta_2\, AVGMARGDIFF_{ij} + \beta_3\, LNAVGPRIZE\mathbf{FEMALE}_j + \beta_4\, AVGMARGDIFF\mathbf{FEMALE}_j + \beta_5\, GENDER + \varepsilon_{ij}$

2) $LNTIMEGAP_{ij} = \beta_0 + \beta_1\, LNAVGPRIZE_j + \beta_2\, AVGMARGDIFF_{ij} + \beta_3 LNAVGPRIZE\mathbf{FEMALE}_j + \beta_4\, AVGMARGDIFF\mathbf{FEMALE}_j + \beta_5\, GENDER + \varepsilon_{ij}$

Combining these models with the key principles of tournament theory, several testable hypotheses emerge:

1) The larger the average prize per prize-eligible place in a given race, the faster the finishing times of its top-finishing participants and the lower the time gap behind the winner;

2) The higher the average marginal returns to finishing in an adjacent place (one higher or one lower), the faster runner i's performance and the lower the time gap behind the winner;

3) The effect of incentive structure on performance is stronger for male runners than for female runners.

The results of an Ordinary Least Squares estimation for the first model are shown in Table 3. From this regression, several initial observations can be made. First, for the whole sample, both incentive variables are significant at at least the 95% level and show a negative relation to runner performance, all other variables held equal. A 10% increase in average prize money per prize-eligible place (LNAVGPRIZE) leads, in this sample, to a 20% decrease in an individual's finishing time (i.e. a faster performance), all else equal. Using the sample means of these variables to get a better sense of the actual figures involved, this coefficient entails that for an increase of $110 in average prize money awarded in a given race, a runner will run approximately 3 minutes faster. Looking at the other incentive variable, the coefficient on average marginal change to finishing in an adjacent place (AVGMARGDIFF) also shows a negative relationship with finishing time, all other variables held equal. A 10% increase in this variable leads, in this sample, to a 0.8% decrease in an individual's finishing time, all else equal. Again using the sample means, this coefficient entails that for an increase of $580 in average marginal difference between adjacent places, a runner will run approximately 71 seconds faster.

When these two incentive variables are interacted with the gender term (0/1 for male/female), they lose much of their statistical significance. While the coefficient on the term interacting average prize with gender shows that this variable is 1.4% stronger for female runners than for male runners (i.e. if coefficient = 0), it does so only at a 90% significance level. Still, this result is notable, as the coefficient is still somewhat significant, and a stronger incentive effect for females would conflict with the trends found in the existing experimental literature (more on this in the next section).   Looking at the other interacted incentive variable, however, the coefficient is significant only at a 72% level, suggesting that males and females respond more or less similarly to marginal differences in prizes.

Table 3: The Determinants of Finishing of Times

Variable	**β**	**T-Statistic**
LNAVGPRIZE	-0.020146	*-3.33****
AVGMARGDIFF	-0.0008211	*-2.3***
LNAVGPRIZE**FEMALE**	-0.0148874	*-1.74**
AVGMARGDIFF**FEMALE**	-0.0005476	*-1.08*
GENDER	0.1825405	*10.38***
CONSTANT	9.058389	*728.69****
* p < .10 ** p < .05 *** p < .01		(n = 216)

The regression results of the second model, regressing the same set of variables on time gap, a measure of race spread or closeness, are shown in Table 4. Of the two incentive variables, only average prize is significant, showing a negative relation with time gap (i.e. a tighter spread). The coefficient on this variable entails that a 10% increase in average prize is related to an enormous 445% decrease in the average time behind the winner for the second-to fifth-place finishers in a given race, all other variables held equal. Using the means of the sample for some perspective, an $1100 increase in the average prize in a given race leads to a decrease in time gap of approximately 19 minutes. In other words, the greater the average prize money per eligible place, the more tightly clustered a marathon race, all else equal. Of the other variables in the model, the only other that is significant is the term interacting average prize with gender. What is different for this term, though, is its positive sign. The coefficient on this variable entails that the effect of average prize on race closeness is 62% stronger in female races than in male races, but that the larger the average prize, the *more* dispersed a marathon's top finishers.

Table 4: The Determinants of Race Closeness

Variable	**β**	**T-Statistic**
LNAVGPRIZE	-0.4451008	*-3.78****
AVGMARGDIFF	-0.0054486	*-0.78*
LNAVGPRIZE**FEMALE**	0.6204099	*3.72****
AVGMARGDIFF**FEMALE**	0.0057423	*0.58*
GENDER	-0.548857	*-1.6*
CONSTANT	5.715676	*23.59****
* p < .10 ** p < .05 *** p < .01		(n = 216)

IX. Commentary

Looking at the regression results of the two models, the testable hypotheses marked out at the start of the previous section are met with mixed results. The first hypothesis is confirmed by the data: the larger the average prize in a given marathon, the faster its prize-eligible runners will perform, and the more tightly bunched those runners will finish. The negative relation between average prize and finishing time is not surprising, though it does have one large potential problem of endogeneity. While larger prizes are correlated with faster times, it remains unclear whether the link is causal, or if it demonstrates a selection effect: that is to say, if prizes with larger purses simply attract faster runners, rather than inducing them to run faster. This selection effect might explain, furthermore, the relation between average prize and time gap. A marathon that attracts fast runners will also attract many of them, creating a deeper, more competitive field more likely to finish close together. It is clear, though, that larger average prizes will help to produce tighter competitions.

In testing the principles of tournament theory in a non-experimental athletic setting, though, what one seeks to target is a potential effort effect: not simply the selection into a lucrative tournament environment, but the differences in effort induced by the features of such an environment. With this end in mind, the average marginal difference variable can be potentially informative, picking up incentive effects not captured by average prize, which is a less subtle, more aggregate measure. In line with Lazear and Rosen's theoretical modeling, the data show that the larger the prize spread between adjacent ranks, the better an individual will perform (that is to say, the more effort an individual will exert), all else equal. Such a finding is logical, as the greater the magnitude of a one-place change, the more effort will be put forth—motivated either by a larger marginal return for a higher place or a larger opportunity cost for finishing in a lower place. One would assume, though, that this overall increase in effort would lead to tighter races, or races with smaller measures of the time gap variable. The data of the sample do not, however, support this assumption, as the coefficient on average marginal difference in the second model is not significant.

The story of gender differences as demonstrated by these results is largely absent, though, as male and female runners seem not to differ significantly in their performance responses to variations in incentive structures. While the variance in the performance variables (time and time gap) for female runners in the sample is much higher than that for males, such variance is not a function of the incentive variables (at any significant level). The relevant literature on gender differences in competitive environments, reviewed previously in this paper, might seem to conflict with this finding. In the series of experiments run in those papers, three general trends emerged. First, female performance was found to be virtually identical under noncompetitive incentive schemes (piece rate) as under competitive (rank-order) schemes; second, effects of intra-gender competition are stronger for males than for females; and third, selection into competitive environments is more likely for males than for females. But while this literature marks out some general trends, it is crucially different from the non-experimental spaces from which my own data comes. The existing literature relies upon a random sample of participants; professional marathoners, the subjects of my sample, are far from a random group. While it is possible, these papers suggest, that real gender differences in behavioral preferences do exist in the general population, viewing subjects who have invested resources in pursuing professional marathoning as similar to a random group is bound to be problematic. What I have attempted to identify in my research is not whether males and females respond differently to competitive environments (this has already been identified), but whether such differences extend to competitive male and female marathon runners. My sample has shown that such differences do not in fact extend to this particular type of contest, or at least not at a level that is statistically significant. Professional male and female marathoners can be seen, then, to be similarly inclined with respect to their performance in response to competitive incentive schemes.

X. Summary and Conclusions

Using a data set of 216 runners from the 12 largest domestic marathons in 2005, I have considered first, the effects of incentive structures on runner performance and marathon competitiveness and second, potential gender differences in these effects. While my models revealed a strong connection between average prize money per prize-eligible place and runner performance (a positive relation, with finishing times decreasing) and race competitiveness (tighter time spreads), considering only this variable was not sufficient, as the impact of selection could not be ruled out. My other incentive variable, the average marginal difference in finishing in an adjacent place, targeted more closely the impact of prize structure on effort. The sample showed that the greater the prize spread between adjacent places, the faster an individual's performance. This is a result that confirms a basic principle of tournament theory from personnel economics—that larger compensation spreads between adjacent ranks induce greater worker effort. My second consideration for the sample was potential gender differences in these incentive effects. Seemingly at odds with much of the existing experimental work on the question, the sample showed only insignificant differences between male and female runners regarding the impact of prize structure on performance. Given the prerequisite of competitiveness likely common to all professional-level marathon runners, gender differences observed among random subjects were not seen to extend to my sample—one that is far from randomly chosen.

While my results, especially those regarding the relation between marginal spreads and effort/performance, demonstrate a useful application of tournament theory to marathon running, they are largely preliminary. Though difficult to find, a larger panel data set would enable fixed effects regressions, a model that could potentially control for the selection effect that marks my first incentive variable. Additionally, a larger sample of marathons might enable considerations on a more comprehensive, race level as opposed to simply on the level of individual competitors. Top marathon runners tend to be active on the marathon circuit for only a few years, so the available data may cater more to considerations regarding the design of an optimal marathon, rather than, more narrowly, to considerations regarding how to optimally incentivize runner performance.

XI. Appendix 1: Selected Prize Distributions

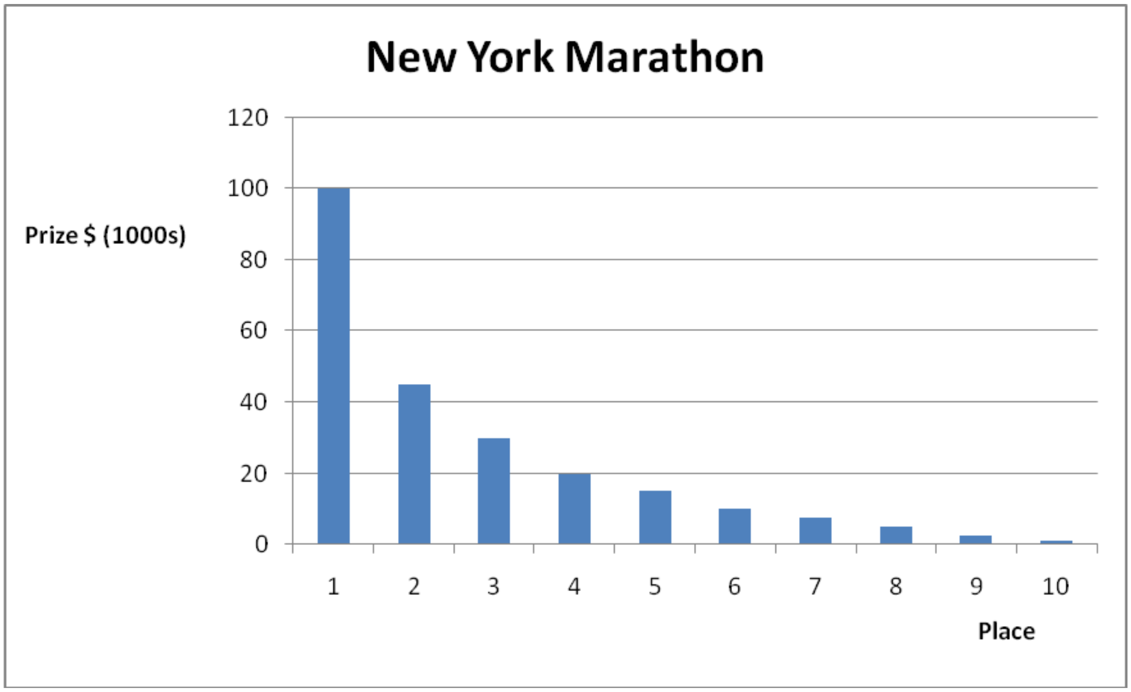

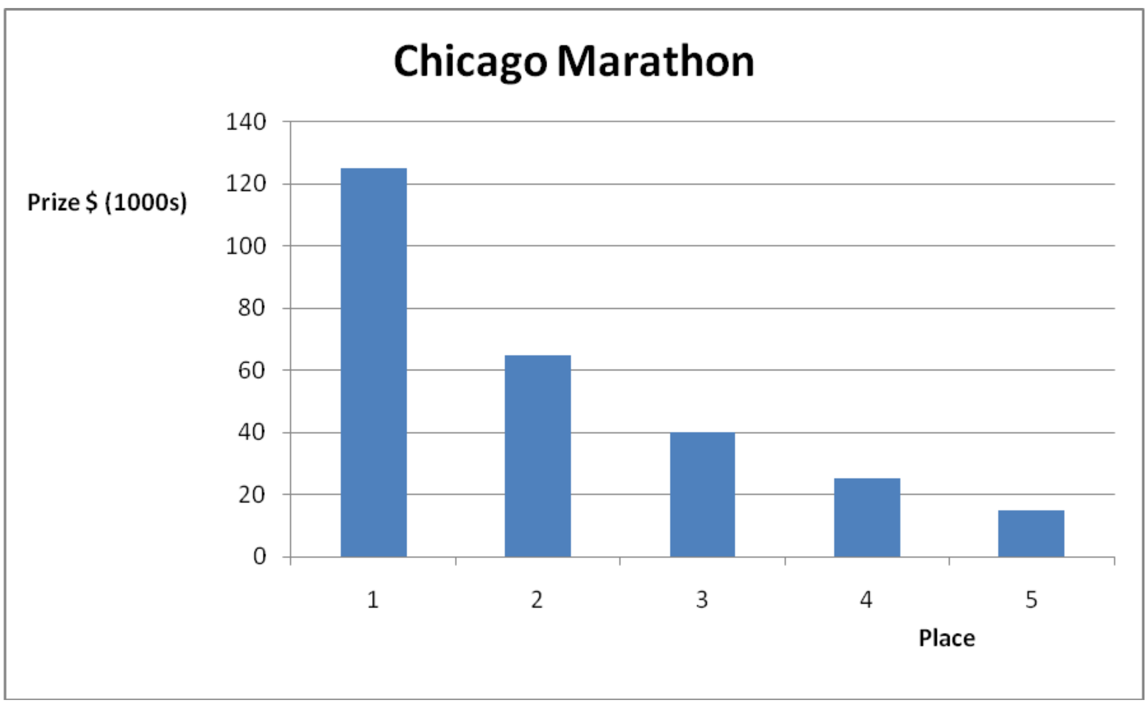

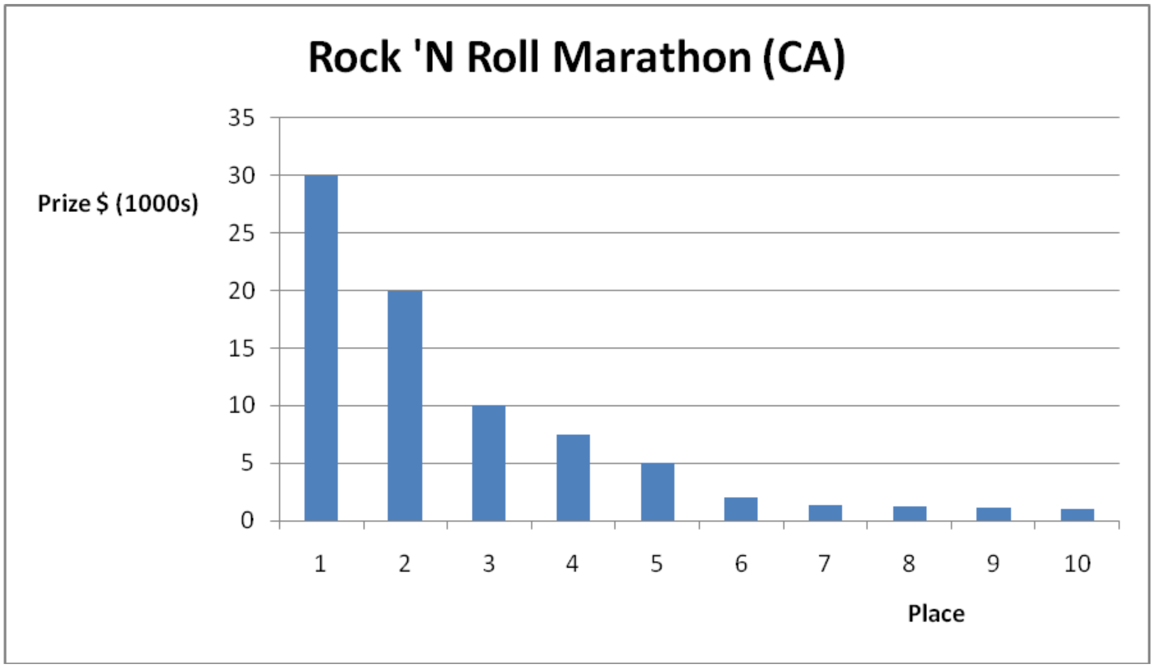
Rock 'N Roll Marathon (CA)
Prize $ (1000s)
35
30
25
20
15
10
5
0
1
2
3
4
5
6
7
8
9
10
Place

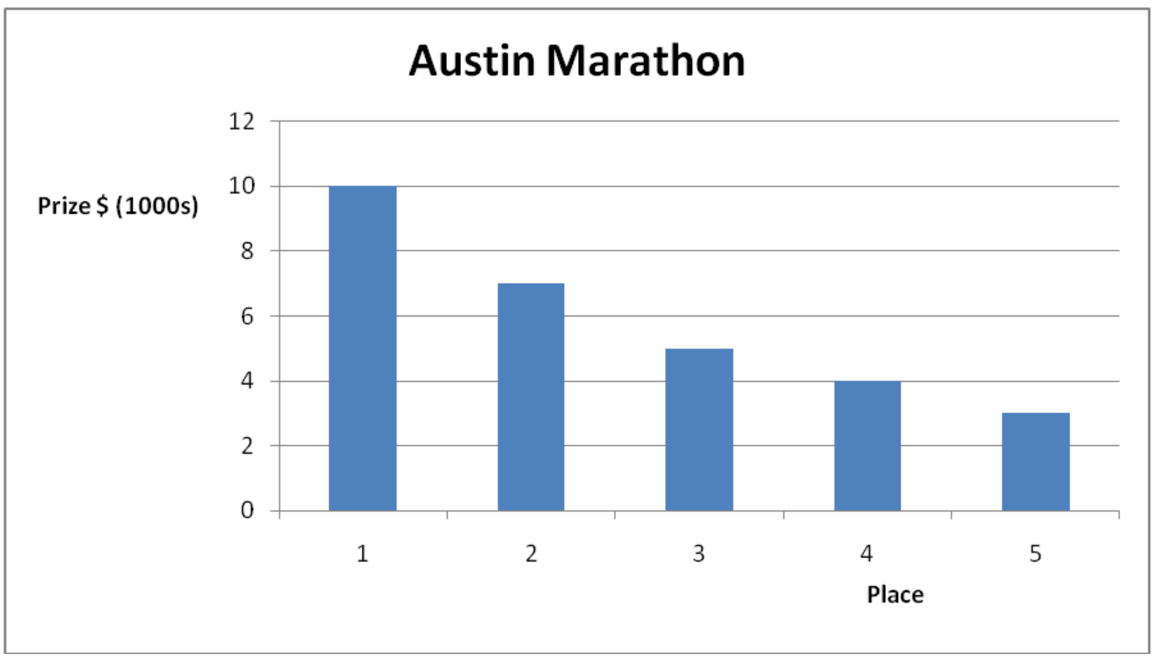
Austin Marathon
Prize $ (1000s)
12
10
8
6
4
2
0
1
2
3
4
5
Place

XII. References

Becker, B.E. and Huselid, M.A. (1992). "The Incentive Effects of Tournament Compensation Systems." *Administrative Science Quarterly*, 37, 336-350.

Black, Sandra and Strahan, Philip (2001). "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry." *The American Economic Review*, 91:4, 814-831.

Demsetz, Harold (1995). *The Economics of the Business Firm: Seven Critical Commentaries*. New York: Cambridge University Press.

Ehrenberg, R.G. and Bognanno, M.L. (1990). "Do tournaments have incentive effects?" *Journal of Political Economy*, 98, 1307-1324.

Frick, Bernd and Prinz, Joachim (2007). "Pay and Performance in Professional Road Racing: The Case of City Marathons." *International Journal of Sport Finance*, 2:1, 25-35.

Gneezy, Uri and Rustichini, Aldo (2002). "Gender and Competition at a Young Age." *American Economic Review*, 94:2, 377-381.

Gneezy, Uri and Niederle, Murial (2003). "Performance in Competitive Environments: Gender Differences." *The Quarterly Journal of Economics*, 118:3, 1049-1074.

Lazear, Edward and Rosen, Sherwin (1981). "Rank-Order Tournaments as Optimum Labor Contracts." *The Journal of Political Economy*, 89:5, 841-864.

Lynch, James G. and Zax, Jeffrey S. (2000) "The Rewards to Running: Prize Structure and Performance in Professional Road Racing." *Journal of Sports Economics*, 1:4, 323-340.

Nierderle, Murial and Vesturlund, Lise (2007). "Do Women Shy Away From Competition? Do Men Compete Too Much?" *The Quarterly Journal of Economics*, 122:3, 1067-1101.

Szymanski, Stefan (2003). "The Economic Design of Sporting Contests." *Journal of Economic Literature*, XLI, 1137-1187.