Lindner, Florian; Dutcher, E. Glenn; Balafoutas, Loukas; Ryvkin, Dmitry; Sutter, Matthias

Conference Paper

Strive to be first and avoid being last: An experiment on relative performance incentives

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2013: Wettbewerbspolitik und Regulierung in einer globalen Wirtschaftsordnung - Session: Pay and Performance, No. C17-V1

Provided in Cooperation with:
Verein für Socialpolitik / German Economic Association

Strive to be first or avoid being last: An experiment on relative performance incentives.[†]

E. Glenn Dutcher[‡,*], Loukas Balafoutas[*], Florian Lindner[*], Dmitry Ryvkin[§] and Matthias Sutter[*,¶]

July 2, 2013

Abstract

In this paper we compare effort provision and efficiency under incentive schemes which motivate agents to compete to be first, avoid being last, or both. We utilize a laboratory experiment to test the predictions generated by optimal contracts in a Lazear-Rosen type model that predicts equal effort and efficiency levels for the three mechanisms. We find that the mechanism incorporating both incentive schemes induces the highest effort, especially in contests of relatively larger size. Avoiding being last is shown to produce the lowest variance of effort, be more effective and, in larger contests, more efficient than competing for the top.

JEL Codes: M52 J33 J24 D24 C90

Keywords: tournament, winner, loser, contract, experiment, learning

[†]Financial support from the Austrian Science Fund (FWF) through project P22772-G11 is gratefully acknowledged. We would also like to acknowledge the useful feedback received from seminars at the University of Texas Dallas, the University of Innsbruck, the University of Linz and MPI - Munich as well as participants at the 2012 Economic Science Association Meetings in New York and Tucson, the French Experimental Economics Association Meeting in Montpellier and the Southern Economics Association in New Orleans.

[‡]Corresponding author: Department of Public Finance, University of Innsbruck, Universitaetsstrasse 15, A-6020 Innsbruck Austria, glenn.dutcher@uibk.ac.at

[*]University of Innsbruck.

[§]Florida State University.

[¶]University of Cologne.

1 Introduction

Managers in organizations have many motivational tools at their disposal. A popular such tool is the use of incentive schemes based on ordinal relative performance evaluations, or rank-order tournaments. A recent Wall Street Journal article states that 60% of Fortune 500 companies currently use some kind of a ranking system for incentive provision.[1] Thus, in these companies employees compete with each other for rewards such as promotions and bonuses. The popularity of such mechanisms is largely due to an inherent structure present in most organizations, where only a limited number of promotions or amount of bonus money exists. With this natural limitation, managers must be selective in whom to give the reward(s) to.[2] The best or highest performing employee(s) will often get the nod, which gives employees the incentive to work harder.

Given the prevalence of tournament-based incentive systems in organizations, it is no wonder that this topic has generated a magnitude of economic research (see surveys in Konrad 2009 and Dechenaux, Kovenock and Sheremeta 2012). However, most of the literature on rank-order tournaments focuses on understanding how participants compete for the top prize(s), while relatively little research has been directed into understanding how incentive schemes motivating participants to avoid being last affect employee effort. For instance, the aforementioned WSJ article also states that when Country Wide had to lay off employees, they first selected those who were ranked the lowest from prior evaluations. Though termination is the most severe form of being ranked last, it need not be the only form. It is often the case that lower ranked employees are demoted, assigned to less desirable tasks, have bonuses withheld, etc.

A clean identification of the incentive effects of tournaments is quite difficult, since data collected in the field usually allow one to observe only outcomes (i.e., total output). This is problematic since in most instances outcomes are a function of luck, noise, ability, endogenous selection as well as effort. Due to the difficulties in isolating the incentive effects, laboratory experiments have often been utilized as a way of giving more control to the researcher. A brief review of these experiments is given in Section 2. At this point, we note that the focus of experimental research has been centered mainly on competition for the top prize(s).

[1] The article was titled "'Rank and Yank' Retains Vocal Fans" which was published on January 31, 2012 and can be accessed via the following link. http://online.wsj.com/article/SB10001424052970203363504577186970064375222.html

[2] As pointed out in Lazear and Rosen (1981), one further reason for the popularity of rank-order tournaments has to do with complications inherent in organizations, which inhibit a manager from forming contracts on effort directly due to common productivity shocks or the difficulties in measuring actual output in a quantifiable way.

This paper aims to complement the existing literature by adding to it an empirical comparison of tournament mechanisms involving competition for the top, competition to avoid the bottom, or both in a principal-agent setting. Using the prominent theory of Lazear and Rosen (1981), we examine the following three tournament mechanisms where agents compete for various prizes: A *winner tournament* is a mechanism where the agents compete to be first, and one top prize is awarded to the agent with the highest output. Likewise, in a *loser tournament* the agents compete to avoid being last, and one bottom prize is given to the agent with the lowest output. Finally, a *winner&loser tournament* is a combination of the two. We first use the theory of Lazear and Rosen (1981) to calculate the optimal principal-agent contracts for each of these mechanisms. The model predicts that all three mechanisms are efficient contracts and that the same efficient levels of effort are exerted in each case – as noted in Lazear and Rosen (1981) and Nalebuff and Stiglitz (1983). We then parameterize the theoretical model and directly test the hypotheses derived from it in a laboratory experiment, in which subjects in the role of employees choose effort levels (tied to a convex cost structure) and compete in one of the three mechanisms defined above.

In our design, not only is employee effort predicted to be the same in all three mechanisms, but the employer's per-employee cost is also the same. Hence, theory predicts no difference between mechanisms. Our experiment allows us to test this prediction, as well as to study the properties of each mechanism in isolation. In addition to these three mechanisms, we also vary the size of the tournament, considering tournaments of size three and of size six. Varying the size of the tournament serves two primary purposes. First, tournaments in organizations vary in size. Understanding how the different mechanisms interact with the size of the tournament is highly relevant to forming generalizable recommendations. Second, varying the number of contestants in the tournament will help us to disentangle the underlying causes of the differences observed between mechanisms, and will also provide robustness to our results.

Our findings show that, in contests of size three, the winner tournament is inferior to the other two in terms of employee effort, while there is no clear distinction between the mechanisms in terms of aggregate efficiency (measured as the combined gains to the principal and the agents). The existence of a top prize in the winner tournament encourages some subjects to choose very high effort, but on the other hand a large number of subjects choose extremely low effort due to the stiff competition for the top prize. By contrast, the mechanism which only includes a single bottom prize practically eliminates the subjects' effort choices in the lowest range, although it also weakens the incentive for subjects to provide high effort. In contests of size six, the winner and loser tournaments

(the ones with a single top or single bottom prize respectively) turn out to produce statistically indistinguishable results. Using a mechanism that combines both a top prize and a bottom prize brings out the best of both mechanisms, and this turns out to be the most efficient mechanism regardless of contest size. Under this scheme, competition away from the bottom reduces the number of subjects who choose low effort, while the competition for the top provides continuous encouragement for some subjects to choose very high effort.

In terms of a direct comparison of winner and loser tournaments, we find that the loser tournament is overall a better motivator than the winner tournement. This leads to the finding that when the tournament size is small, the loser tournament approaches the outcome attained in the winner&loser tournament. This result is largely in line with the predictions generated by basic learning theory. Note, in smaller contests the probability of being the lowest performing employee is higher and thus being the lowest performer will be experienced more often. Over time, this effect leads to similar outcomes for loser tournaments and winner&loser tournements when the number of contestants is small. The strong learning effect we find also implies that a manager using a rank-order pay scheme must use the mechanism regularly as the increases in work effort will decline if the incentives are absent.

2 Brief review of the related literature

Extensive theoretical work has been undertaken to understand rank-order tournaments in an organizational setting. This literature has mainly focused on winner tournaments (see reviews by McLaughlin 1988; Lazear 1995 and Konrad 2009).[3] Loser tournaments were first mentioned by Mirrlees (1975), while Nalebuff and Stiglitz (1983) prove the equivalence of the two schemes in the more general symmetric setting. Loser tournaments have more recently started to attract theoretical examinations by authors looking at heterogeneity in some aspect (e.g. Kräkel 2000, Gürtler and Kräkel 2011, Moldovanu, Sela and Shi 2012 and Balafoutas et al. 2012) or risk-aversion (Akerlof and Holden 2012).

Many experiments have also been conducted on tournaments (see a review by Dechenaux, Kovenock and Sheremeta 2012). The first study to examine rank-order tournaments, conducted by Bull, Schotter and Weigelt (1987), found that tournament and piece-rate pay

[3]Our main interest is in the static principal-agent models of tournaments à la Lazear and Rosen (1981). Dynamic tournaments which involve sequential elimination of employees has also been explored (see, e.g., Rosen 1986, O'Flaherety and Siow 1995, Gradstein and Konrad 1999, Ryvkin and Ortmann 2008, Casas-Arce and Martinez-Jerez 2009, Sunde 2009, Höchtl et al. 2011). Even though the elimination mechanism can be thought of as a loser tournament, the focus of these papers is not on this aspect, but on the effect elimination has on the remaining agents.

schemes generated the same mean effort, though the tournament pay scheme induced a higher variance in effort. With the very basic tenant of tournament theory established, subsequent papers focused on other topics such as affirmative action (Schotter and Weigelt 1992), tournament size and prize structure (Harbring and Irlenbusch 2003, Orrison, Schotter and Weigelt 2004), sabotage (Harbring and Irlenbusch 2008, Falk, Fehr and Huffman 2008, Carpenter, Matthews and Schirm 2010, Harbring and Irlenbusch 2012), selection (Camerer and Lovallo 1999, Eriksson, Teyssier and Villeval 2009, Cason, Masters and Sheremeta 2010, Müller and Schotter 2010), dynamic tournaments (Sheremeta 2010), and gender effects (Gneezy, Niederle and Rustichini 2003) among others. The most relevant for us are those that vary the number of contestants in a rank-order tournament as well as the number of winner prizes.

The two prior studies that are most relevant to our study (Orrison, Schotter and Weigelt 2004, Harbring and Irlenbusch 2008) vary the fraction of winner and loser prizes in different sized tournaments. The overall finding in these papers is that the highest effort is observed when there is an equal distribution of winner and loser prizes. This is counter to Harbring and Irlenbusch (2003) who show that in a capped all-pay auction setting, effort increases with the number of winner prizes. Concerning contest size, Orrison, Schotter and Weigelt (2004) find that effort does not depend on the size of the contest. Harbring and Irlenbusch (2008) also find that there is no discernible trend that relates effort to the size of the contest. Using non-uniform distributions of noise, List et al. (2010) support this result for risk-neutral subjects, but find that risk-averse subjects' effort declines as the number of contestants increases.

Although these papers have treatments that are similar in concept to our design, key differences remain between them and our study due to the disparate research questions. First, the above studies only implement two levels of prizes. In our design, we will implement a treatment with a top, middle and bottom prize (the winner&loser tournament). This is an important variation for several reasons. Since most organizations typically use incentive schemes which has consequences for the best and the worst workers, it is important to understand how effort and efficiency may differ under this very common pay scheme which uses three distinct levels of prizes. We will show that including three distinct prizes is essential since managers who choose this pay scheme can typically expect higher effort. More importantly, since the above papers were not intended to test the differences between competition for the top or away from the bottom, they do not use optimal contracts – a feature which is necessary for a meaningful efficiency comparison and a potential cause of the partly contradictory results of the papers mentioned above. More specifically, optimal contracts are important since they fix the amount of money

paid per worker to be equal in every setting while holding constant the predicted effort. Thus, variations that we observe in our setting can be attributed to the underlying behavioral response to winner and/or loser tournaments, and efficiency concerns can be cleanly examined.

3 The model

3.1 Three tournament mechanisms

There are $n \geq 2$ identical risk-neutral agents indexed by $i = 1, \ldots, n$.[4] Each agent participates in the tournament by exerting effort $e_i \geq 0$. The cost of effort e_i to agent i is $cg(e_i)$, where $c > 0$ is the agents' homogeneous cost parameter, and function $g(\cdot)$ is strictly increasing and strictly convex.

As in Lazear and Rosen (1981), agent i's output is $y_i = e_i + u_i$, where u_i is a zero-mean idiosyncratic random shock. It is assumed that shocks $u_1, \ldots, u_n$ are i.i.d. across agents and drawn from the distribution with support $[u_l, u_h]$, pdf $f(u)$ and cdf $F(u)$.

In tournament mechanisms, agents are evaluated on the basis of their relative performance. Effort is not observable and cannot be used for contracting. Moreover, cardinal output is also not observable. Contracts are based on ordinal comparisons of agents' output levels. We consider three tournament mechanisms: winner tournaments, loser tournaments, and winner&loser tournaments.

Winner tournament is defined as a tournament in which the winner, i.e., the agent with the highest output, receives prize V_1, and all other agents receive prize V_2, with $V_2 < V_1$. The probability of agent i winning the tournament, $p^{(i)}$, is[5]

$$p^{(i)}(\mathbf{e}) = \int \left[\prod_{j \neq i} F(t + e_i - e_j) \right] f(t) dt. \tag{1}$$

Here, $\mathbf{e} = (e_1, \ldots, e_n)$ is the vector of all agents' effort levels. Agent i's expected payoff is

$$\pi_i(\mathbf{e}) = V_2 + (V_1 - V_2) p^{(i)}(\mathbf{e}) - cg(e_i).$$

Assuming all agents participate in the tournament with positive efforts,[6] the vector of

[4] See Akerlof and Holden (2012) for a theoretical examination of how risk-aversion affects the predictions.

[5] Integration everywhere is over the support $[u_l, u_h]$ of the distribution of noise.

[6] Note that participation in the tournament with zero effort is not equivalent to nonparticipation. Because of noise, the probability of winning the tournament with zero effort is still positive.

equilibrium effort levels, $\mathbf{e}^* = (e_1^*, \ldots, e_n^*)$, solves the system of first-order conditions[7]

$$(V_1 - V_2) \sum_{j \neq i} \int \left[\prod_{k \neq i,j} F(t + e_i - e_k) \right] f(t + e_i - e_j) f(t) dt = cg'(e_i), \quad i = 1, \ldots, n. \quad (2)$$

Loser tournament is defined as a tournament in which the agent with the lowest output receives prize W_2, and all other agents receive W_1, with $W_1 > W_2$. The probability of agent i having the lowest output, $q^{(i)}$, is

$$q^{(i)}(\mathbf{e}) = \int \left[\prod_{j \neq i} (1 - F(t + e_i - e_j)) \right] f(t) dt. \quad (3)$$

Agent i's expected payoff is

$$\pi_i(\mathbf{e}) = W_1 - (W_1 - W_2) q^{(i)}(\mathbf{e}) - cg(e_i).$$

Assuming full participation with positive efforts, the vector of equilibrium effort levels $\tilde{\mathbf{e}}^* = (\tilde{e}_1^*, \ldots, \tilde{e}_n^*)$, solves the system of first-order conditions

$$(W_1 - W_2) \sum_{j \neq i} \int \left[\prod_{k \neq i,j} (1 - F(t + e_i - e_k)) \right] f(t + e_i - e_j) f(t) dt = cg'(e_i). \quad i = 1, \ldots, n. \quad (4)$$

Winner&loser tournament is defined as a tournament in which both the winner and loser elements are combined: the agent with the highest output receives a prize S_1, the agent with the lowest output receives a prize S_3, and all other agents receive S_2, with $S_1 > S_2 > S_3$. Agent i's expected payoff in this setting can be written as

$$\pi_i(\mathbf{e}) = S_2 + (S_1 - S_2) p^{(i)}(\mathbf{e}) - (S_2 - S_3) q^{(i)}(\mathbf{e}) - cg(e_i).$$

Assuming full participation with positive efforts, the vector of equilibrium effort levels $\hat{\mathbf{e}}^* = (\hat{e}_1^*, \ldots, \hat{e}_n^*)$, solves the following system of first-order conditions that will contain a sum of the terms from the left-hand sides of Eqs. (2) and (4), multiplied by prize differentials $S_1 - S_2$ and $S_2 - S_3$, respectively:

[7] The first-order conditions derived in this section are necessary but not sufficient for the equilibrium to exist. We discuss the corresponding restrictions on parameters in Section 4.3.

$$(S_1 - S_2) \sum_{j \neq i} \int \left[\prod_{k \neq i,j} F(t + e_i - e_k) \right] f(t + e_i - e_j) f(t) dt$$
$$+ (S_2 - S_3) \sum_{j \neq i} \int \left[\prod_{k \neq i,j} (1 - F(t + e_i - e_k)) \right] f(t + e_i - e_j) f(t) dt$$
$$= cg'(e_i). \; i = 1, \ldots, n. \quad (5)$$

There is a risk-neutral principal, whose expected payoff (firm's profit) is defined as the difference between aggregate effort and total prize payments: $\Pi = \sum_i e_i - V_1 - (n-1)V_2$ for winner tournaments, $\Pi = \sum_i e_i - (n-1)W_1 - W_2$ for loser tournaments, and $\Pi = \sum_i e_i - S_1 - (n-2)S_2 - S_3$ for winner&loser tournaments. In the derivation of optimal contracts, we follow the approach of Lazear and Rosen (1981) and assume that the principal operates in a (buyer-side) competitive labor market under the zero-profit condition $\Pi = 0$.

3.2 Symmetric optimal contracts

We restrict the analysis to the symmetric case in which all agents exert the same effort in equilibrium. For winner tournaments, the first-order condition, Eq. (2), for agents' equilibrium effort, $\bar{e}$, is

$$(V_1 - V_2)\alpha_n = cg'(\bar{e}), \qquad \alpha_n = (n-1) \int F(t)^{n-2} f(t)^2 dt. \quad (6)$$

If Eq. (6) has a solution $\bar{e} > 0$ (which is the case provided $(V_1 - V_2)\alpha_n > cg'(0)$), it is unique. The principal's profit is $\bar{\Pi} = n\bar{e} - V_1 - (n-1)V_2$, which implies, under the zero-profit condition, $\bar{e} = V_1/n + (n-1)V_2/n$, and the agents' expected payoffs are $\bar{\pi} = \bar{e} - cg(\bar{e})$. As in Lazear and Rosen (1981), we assume that the principal chooses prizes V_1 and V_2 to maximize $\bar{\pi}$, implying $(1 - cg'(\bar{e}))\partial\bar{e}/\partial V_k = 0$, $k = 1, 2$. This gives the following system of equations for the optimal contract:

$$(V_1 - V_2)\alpha_n = cg'(\bar{e}), \qquad n\bar{e} = V_1 + (n-1)V_2, \qquad cg'(\bar{e}) = 1. \quad (7)$$

For loser tournaments, the symmetric first-order condition for agents' equilibrium effort, $\tilde{e}$, is

$$(W_1 - W_2)\tilde{\alpha}_n = cg'(\tilde{e}), \qquad \tilde{\alpha}_n = (n-1) \int [1 - F(t)]^{n-2} f(t)^2 dt. \quad (8)$$

If Eq. (8) has a solution $\tilde{e} > 0$ (which is the case provided $(W_1 - W_2)\tilde{\alpha}_n > cg'(0)$), it is unique. The principal's profit is $\tilde{\Pi} = n\tilde{e} - (n-1)W_1 - W_2$. Similar to winner contests, the zero-profit condition and maximization of agents' expected payoff leads to the system of equations for optimal contracts (W_1, W_2):

$$(W_1 - W_2)\tilde{\alpha}_n = cg'(\tilde{e}), \qquad n\tilde{e} = (n-1)W_1 + W_2, \qquad cg'(\tilde{e}) = 1. \tag{9}$$

For winner&loser tournaments, the symmetric first-order condition for agents' equilibrium effort, $\hat{e}$, is

$$(S_1 - S_2)\alpha_n + (S_2 - S_3)\tilde{\alpha}_n = cg'(\hat{e}). \tag{10}$$

The principal's profit is $\hat{\Pi} = n\hat{e} - S_1 - (n-2)S_2 - S_3$. Similar to the other two incentive schemes, the zero-profit condition and maximization of agents' expected payoff leads to the system of equations for optimal contract (S_1, S_2, S_3):

$$(S_1 - S_2)\alpha_n + (S_2 - S_3)\tilde{\alpha}_n = cg'(\hat{e}), \qquad n\hat{e} = S_1 + (n-2)S_2 + S_3, \qquad cg'(\hat{e}) = 1. \tag{11}$$

By comparing Eqs. (7), (9) and (11), it is seen that the equilibrium effort is the same under the optimal contracts for all three mechanisms: $\bar{e} = \tilde{e} = \hat{e}$. All three optimal contracts are socially efficient. Individual effort $\bar{e}$ is determined by condition $cg'(\bar{e}) = 1$, and the optimal prizes can be expressed in terms of $\bar{e}$.

For winner tournaments, the optimal contract is

$$\bar{V}_1 = \bar{e} + \frac{n-1}{n\alpha_n}, \qquad \bar{V}_2 = \bar{e} - \frac{1}{n\alpha_n}. \tag{12}$$

For loser tournaments, the optimal contract is

$$\bar{W}_1 = \bar{e} + \frac{1}{n\tilde{\alpha}_n}, \qquad \bar{W}_2 = \bar{e} - \frac{n-1}{n\tilde{\alpha}_n}. \tag{13}$$

For winner&loser tournaments, the number of independent equations for the optimal contract is the same as for the other two mechanisms, but there are three prizes to be determined. This is a manifestation of the more general result, mentioned by Lazear and Rosen (1981), that a tournament mechanism involving any number of prizes between 2 and n can be implemented with only two distinct prizes under symmetry and risk neutrality. Any S_2 between S_3 and S_1 can be implemented as the intermediate prize. For convenience,

we choose $S_2 = (S_1 + S_3)/2$. Equations in (11) then give the following optimal contract:

$$\bar{S}_1 = \bar{e} + \frac{1}{2\hat{\alpha}_n}, \qquad \bar{S}_2 = \bar{e}, \qquad \bar{S}_3 = \bar{e} - \frac{1}{2\hat{\alpha}_n}; \qquad \hat{\alpha}_n = \frac{\alpha_n + \tilde{\alpha}_n}{2}. \tag{14}$$

4 Experimental design and hypotheses

4.1 Basics

Using ORSEE (Greiner 2004), we recruited 216 subjects in total. All experiments were conducted at the University of Innsbruck using z-tree (Fischbacher 2007). Subjects earned on average €12.25 for an experiment that lasted about an hour. Subjects were recruited from the standard subject pool at the University. In our sample, 57% of them were female. Subjects participated in one of 12 sessions, where each session comprised 18 subjects.

4.2 Treatments

The experiment follows a 3×2 between-subjects design covering three tournament mechanisms: winner (WIN), loser (LOS) and winner&loser (W&L); and two group sizes: $n = 3$ and $n = 6$. The resulting six treatments will be referred to as WIN3, WIN6, LOS3, LOS6, W&L3 and W&L6. The procedure used for all 6 treatments was the same. Once subjects were seated in the lab, they were handed instructions which were read aloud to ensure common knowledge of the rules of the game.[8] The first part of the experiment consisted of 20 rounds, and in each round subjects were randomly and anonymously matched in a group with other participants in the session.[9] Each round the subjects participated in a chosen effort task, during which they were asked to choose a number between 1 and 100. Each number had a cost associated with it, with a convex structure.[10]

Following the theory, a random number was added to (or subtracted from) each agent's chosen number, resulting in their "total number."[11] Their total number was then compared to the total number of the other agents in their group to determine their rank within the group. In WIN3 and WIN6, the agent with the highest rank received the top prize while all others received the bottom prize. In LOS3 and LOS6, the agent with the lowest rank received the bottom prize while all others received the top prize. In W&L3 and W&L6

[8] Sample instructions for WIN6 are given in the Appendix.

[9] Random rematching was implemented to reduce reputation effects. Since our main interest lies in understanding how people compete to avoid being last compared to how they compete to be first, including reputation in this design would cloud the main interest of the paper.

[10] At the end of the instruction sheet there was a table showing the cost associated with each number. The cost table is given in the appendix.

[11] The random number was chosen randomly and independently for every agent in each round.

the agent with the highest rank received the top prize, the agent with the lowest rank received the bottom prize, and all others received the middle prize. The total payoff in a period was calculated by subtracting the cost of the chosen number from the prize gained. Effort costs and prizes were denominated in tokens, where 2000 tokens equalled €1. Once all subjects in the session had chosen their number (effort), they were informed of their random number, their total number and whether their total number was the highest (in WIN3, WIN6, W&L3 and W&L6) or the lowest (in LOS3, LOS6, W&L3 and W&L6). Additionally, they were informed of their payment for that round if it would be randomly selected for payment. Four rounds were chosen at random for payment at the conclusion of the experiment.

Before receiving feedback in the last round, subjects were given a Holt-Laury risk aversion task (Holt and Laury 2002), as it is important to control for risk attitudes in our setting given that our theory assumes risk-neutral agents. Subjects were also administered a loss aversion task (Gächter, Johnson and Herrmann 2010), which is meant to control for preferences in the loss domain. The experiment concluded with a short demographic questionnaire.

4.3 Calibration of parameters

In the experiment, we use the uniform distribution of noise on the interval $[-b, b]$. In this case, $\alpha_n = \tilde{\alpha}_n = \hat{\alpha}_n = 1/(2b)$. The optimal contracts for the winner, loser and winner&loser tournaments are

$$\bar{V}_1 = \bar{e} + \frac{2b(n-1)}{n}, \qquad \bar{V}_2 = \bar{e} - \frac{2b}{n};$$

$$\bar{W}_1 = \bar{e} + \frac{2b}{n}, \qquad \bar{W}_2 = \bar{e} - \frac{2b(n-1)}{n};$$

$$\bar{S}_1 = \bar{e} + b, \qquad \bar{S}_2 = \bar{e}, \qquad \bar{S}_3 = \bar{e} - b.$$

For the cost function of effort, we use $g(e) = (A-e)^{-r} - A^{-r}$, with $A, r > 0$. The optimal effort $\bar{e}$ solves $cr(A-e)^{-r-1} = 1$, which gives $\bar{e} = A - (rc)^{1/(r+1)}$. We impose the following three restrictions on parameters.

(a) An agent's payoff $\bar{\pi}$ must be positive, which gives

$$\bar{\pi} = A - (rc)^{1/(r+1)} - c[(cr)^{-r/(r+1)} - A^{-r}] = A + cA^{-r} - (rc)^{1/(r+1)}(1 + r^{-1}) > 0$$

(b) The lowest prize across all the mechanisms, W_2, must be non-negative, which gives

$$A - (rc)^{1/(r+1)} - \frac{2b(n-1)}{n} \geq 0.$$

(c) One more restriction on parameters is imposed by the existence of the symmetric equilibrium in all treatments. This issue is discussed in detail by Akerlof and Holden (2012) and amounts to verification of the second-order condition for the agent's payoff function at the symmetric equilibrium effort $\bar{e}$. Let w_r denote the prize the r^{th} ranked agent receives, and let $\phi^{(r)}(e, \bar{e})$ denote the probability for an agent exerting effort e to be ranked r given that all other agents exert effort $\bar{e}$. The second-order condition for the agent then takes the form

$$\sum_r \gamma_r w_r - \bar{c} g''(\bar{e}) < 0, \tag{15}$$

where $\gamma_r = [\partial^2 \phi^{(r)}(e, \bar{e})/\partial e^2]_{e=\bar{e}}$. The expression for γ_r is provided in Akerlof and Holden (2012) for an arbitrary distribution of noise. It is shown that $\sum_{r=1}^{n} \gamma_r = 0$ and, if the distribution of noise is symmetric, $\gamma_r = \gamma_{n-r+1}$. For winner tournaments, the prize structure is such that $w_1 = \bar{V}_1$ and $w_r = \bar{V}_2$ for $r \geq 2$. Equation (15) then gives

$$\gamma_1(\bar{V}_1 - \bar{V}_2) - \bar{c} g''(\bar{e}) < 0.$$

For loser tournaments, we have $w_r = \bar{W}_1$ for $r \leq n-1$ and $w_n = \bar{W}_2$. This leads to the second-order condition

$$-\gamma_1(\bar{W}_1 - \bar{W}_2) - \bar{c} g''(\bar{e}) < 0.$$

Finally, for winner&loser we have $w_1 = \bar{S}_1$, $w_n = \bar{S}_3$ and $w_r = \bar{S}_2$ for $2 \leq r \leq n-1$. This gives

$$\gamma_1(\bar{S}_1 - 2\bar{S}_2 + \bar{S}_3) - \bar{c} g''(\bar{e}) < 0.$$

Using the equation for γ_r from the proof of Lemma 2 by Akerlof and Holden (2012), for the uniform distribution of noise on $[-b, b]$, we obtain $\gamma_1 = (n-1)/(4b^2)$.[12] This implies that the second-order conditions for the loser and winner&loser tournaments hold automatically as long as $g(\cdot)$ is strictly convex, and the only restriction is imposed by the second-order condition for the winner tournament. The latter gives

$$(rc)^{1/(r+1)} < \frac{2b(r+1)}{n-1}.$$

[12] The complete expression is $\gamma_1 = (n-1)(n-2)\int_{-b}^{b} F(t)^{n-3} f(t)^3 dt$, where $f(t) = 1/(2b)$ and $F(t) = (b+t)/(2b)$ for the uniform distribution.

Mechanism	n	b	c	A	r	Prizes	$\bar{e}$	$\bar{\pi}$
Winner	3	44	3074	106	1.4	$V_1 = 132, V_2 = 44$	73.33	54.49
Loser	3	44	3074	106	1.4	$W_1 = 102.67, W_2 = 14.67$	73.33	54.49
Winner&Loser	3	44	3074	106	1.4	$S_1 = 117.33, S_2 = 73.33, S_3 = 29.33$	73.33	54.49
Winner	6	44	3074	106	1.4	$V_1 = 146.67, V_2 = 58.67$	73.33	54.49
Loser	6	44	3074	106	1.4	$W_1 = 88, W_2 = 0$	73.33	54.49
Winner&Loser	6	44	3074	106	1.4	$S_1 = 117.33, S_2 = 73.33, S_3 = 29.33$	73.33	54.49

Table 1: Treatments and parameters of the experiment.

Table 1 shows parameters of the experiment that satisfy the constraints with the above cost function.[13]

4.4 Hypotheses

Based on the above theory and parameterization, we can formulate the following main hypotheses, which our experiment aims to test:

Hypothesis 1: Subjects' effort is equal in all treatments.

Hypothesis 1 makes predictions about the comparison between treatments, without the numerical restriction that effort be equal to 73 (as equilibrium predicts). Notice that this prediction about equality of efforts does not depend upon the size of the tournament or the mechanism employed.

Hypothesis 2: All treatments are equally efficient.

Efficiency is defined as the combined gains to the principal and the agents. Hypothesis 2 also follows directly from the theory and parameterization. By design, prizes, expected effort and effort costs are the same in each treatment; therefore, efficiency can also be expected to be the same.

5 Results

5.1 Effort

We begin our analysis with an overview of the data. Figure 1 displays the average effort in each period by treatment, in contests of size three (panel a) and size six (panel b). There are 36 observations per treatment per period, so each point in the figure represents the average of these 36 observations. This means that our sample consists of 4,320

[13] There are minor discrepancies in Table 1 due to rounding. In the experiment, all the prizes and costs have been multiplied by 100 to avoid the decimals.

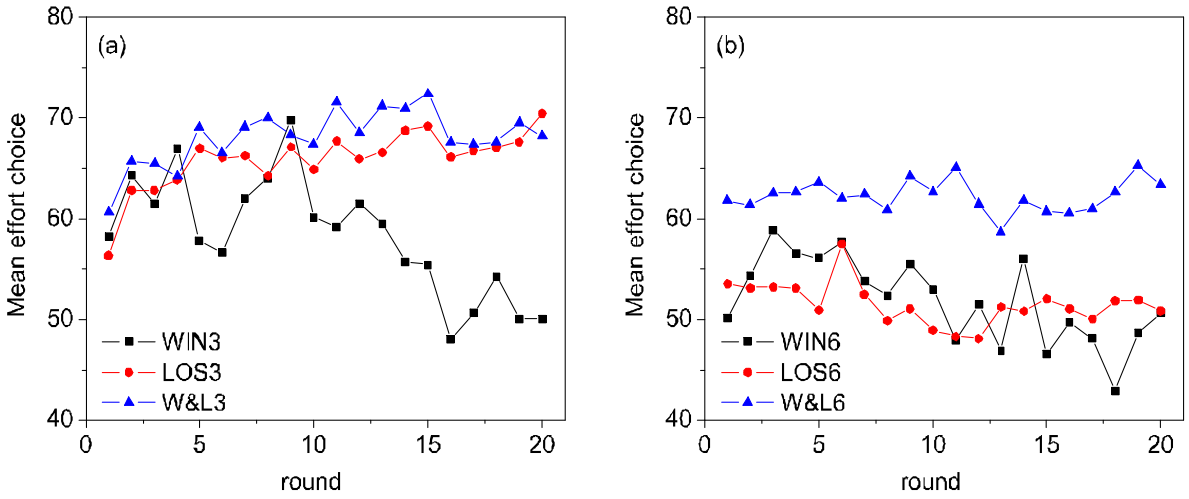

Figure 1: Mean effort over time. Panel (a): Effort in contests of size three. Panel (b): Effort in contests of size six.

Size of contest	WIN	LOS	$p-value$ WIN vs. LOS	W&L	$p-value$ W&L vs. WIN	$p-value$ W&L vs. LOS
$n=3$	58.28	65.88	0.03	68.04	<0.01	0.46
$n=6$	51.86	51.49	0.40	62.25	0.03	<0.01
$p-value$ n=3 vs. n=6	0.28	<0.01		0.06		

Table 2: Average effort by treatment. Reported p-values are from a Mann-Whitney test.

observations in total. Looking first at panel (a), we can see that the average effort in WIN3 is much lower than that in LOS3 or W&L3, especially from round 10 onwards, while similar average effort is observed in LOS3 and W&L3. The picture is to a certain extent different for contests of size six: panel (b) shows that the average effort in W&L6 is higher than that in LOS6 or WIN6, while the winner-only and loser-only contests lead to similar effort levels throughout the course of the experiment. Notice that the mean effort in all six treatments is below the predicted value of 73.[14]

Table 2 shows the mean effort for each treatment and the results of pair-wise comparisons using Mann-Whitney tests. Because each subject played the game for 20 rounds, we treat each subject's average effort choice as one observation, resulting in 36 observations per treatment. The results in Table 2 confirm what is already seen in Figure 1. More specifically, in contests of size three there is no significant difference between LOS3 and W&L3, but effort in WIN3 is significantly lower than in either LOS3 or W&L3. In the

[14]The hypothesis of average effort being equal to 73 is rejected in all treatments ($p<0.01$, two-sided $t-test$ with standard errors clustered by subject).

Effort	**(1)** n = 3	**(2)** n = 6	**(3)** Pooled
Constant	40.56*	60.74***	56.99***
	(22.65)	(22.78)	(16.44)
LOS	7.77**	−0.53	7.49**
	(3.40)	(4.06)	(3.39)
W&L	9.59***	10.73***	9.21***
	(3.32)	(4.06)	(3.31)
Age	0.31	0.35	0.29
	(0.48)	(0.58)	(0.38)
Size Six			−7.14
			(4.58)
Size Six×LOS			−8.07
			(5.30)
Size Six×W&L			1.59
			(5.31)
Risk Aversion	0.76	−0.69	−0.08
	(0.84)	(0.65)	(0.51)
Loss Aversion	0.21	0.05	0.20
	(0.85)	(1.03)	(0.68)
Female	−1.44	−3.52	−2.76
	(2.73)	(3.28)	(2.16)
Round	−0.02	−0.21	−0.11
	(0.13)	(0.13)	(0.09)
# of observations	2160	2160	4320
# of clusters	108	108	216
R-squared	0.04	0.05	0.07

Table 3: Individual random effects panel regressions on effort. Robust standard errors (in parentheses) are clustered at the individual level. Three and two stars represent significance at the one and five percent levels respectively.

contests of size six, no difference is detected between WIN6 and LOS6, but there is statistical support for a significantly higher effort in the W&L6 treatment than in the other two treatments. Finally, when we compare the same mechanism by contest size, average effort in contests of size three is found to be higher in the LOS and W&L treatments, compared to the same treatments in contests of size six. However, there is no size effect in the WIN treatment.

Table 3 reports the results of individual random effects regressions. The dependent variable in our regression is individual effort, while the main explanatory variables are dummy variables for the LOS and the W&L treatments (which implies that the reference group for the regressions is the WIN treatment). Additionally, controls for gender, age, round, a subject's risk preferences and the degree of aversion to losses are also included. To correct for the fact that there are 20 observations per subject, standard errors are clustered at the individual level. Column (1) examines contests of size three, column (2) contests of size six, and column (3) combines contests of both sizes, adding an additional dummy variable for contests of size six as well as interactions of size with LOS and W&L.

The model in column (1) from Table 3 confirms that in contests of size three, effort in LOS3 and in W&L3 contests is higher than in the baseline WIN3 contest. A Wald test also confirms that there is no difference between effort in LOS3 and W&L3 ($p = 0.47$). Column (2) shows that, in contests of size six, there is a sizeable difference between WIN6 and W&L6, as well as a difference between LOS6 and W&L6 ($p < .001$); hence, the combination of a separate top and bottom prize gives rise to the highest output. No difference is detected between WIN6 and LOS6. The results from columns (1) and (2) are thus consistent with the non-parametric tests reported above, while column (3) also leads to the same main results using the entire sample.[15] Column (3) also shows that there is no difference in the treatment effects based on the size of the contest, something which is evidenced by the insignificant interaction terms of contest size with the mechanism dummies. As for the effect of size per se, joint significance tests reveal that average effort in contests of size six is lower than in contests of size three in the cases of LOS (Wald test, $p < 0.01$) and W&L ($p < .001$). Thus, after accounting for the level effects of the contest size, the basic mechanism driving behavior is the same in both contests of size three and six. This analysis leads to our first three results.

Result 1: In contests of size three, there is no difference in effort between the LOS and W&L treatments, but effort in both of these treatments is higher than in the WIN treatment.

Result 2: In contests of size six, effort is highest in the W&L treatment while there is no difference in effort between the WIN and LOS treatment.

Result 3: Effort in LOS and W&L contests of size six is lower than in the same contests of size three. Once the level effect from contests of size three is controlled for however, there is no overall difference between the mechanisms employed depending on contest size.

The first three results are arrived at by examining mean effort by treatment. As evidenced by Figure 1, this may not capture all important aspects of behavior, so we proceed with some more nuanced analysis regarding the distribution of effort choices, shown in Figure 2. This Figure presents the frequencies of chosen effort numbers falling into the intervals 1-10, 11-20, ..., 91-100 for all 20 rounds in contests of size three (panel a) and size six (panel b). Focusing first on panel (b), we see that even though Table 2 shows no difference in the means between treatments WIN6 and LOS6, the way that these means are arrived at is quite different. Looking first at the WIN treatment, we see

[15] Using Wald tests, we confirm that in contests of size six there is no significant difference between LOS6 and WIN6 ($p = 0.89$), but a difference remains between WIN6 and W&L6 ($p = 0.01$).

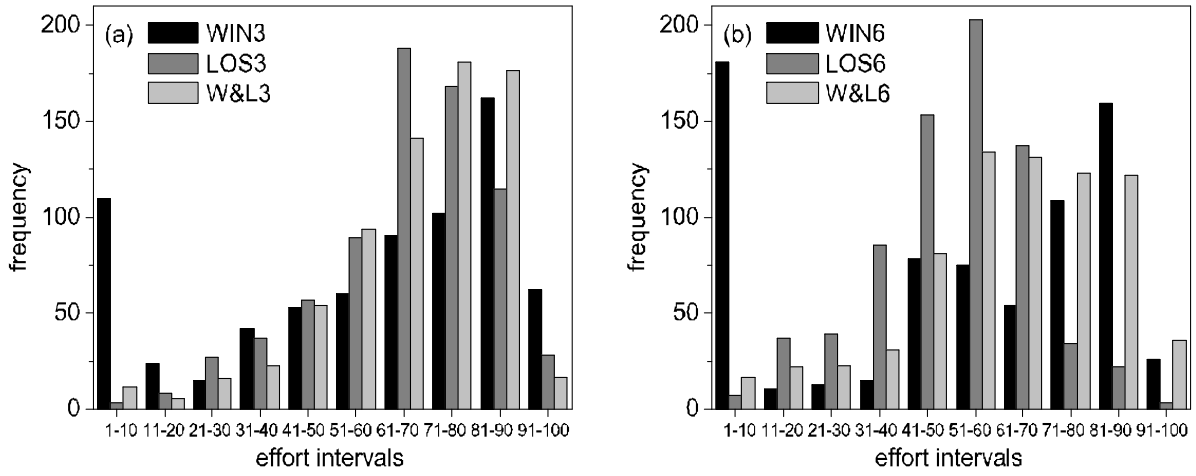

Figure 2: Categories of effort deciles for all 20 periods. Panel (a) is tournaments of size three and Panel (b) is tournaments of size six.

that the distribution of numbers chosen appears bimodal. There are many subjects who contribute 1-10 (the lowest category) and many who contribute 81-90 (the penultimate category). Thus, in line with previous studies (see Dechenaux, Kovenock and Sheremeta 2012), the variance of effort is quite high in the case of winner tournaments, as subjects either compete too much or not enough. In the LOS treatment, the opposite is true. There are very few subjects who choose effort at the lower or upper end, with the majority concentrated at 51-60, close to the mean of 51.5. So the possibility of being last drives up the lower efforts, while the lack of a top prize almost eliminates the higher ones. The W&L mechanism combines both motivations, resulting in a higher overall mean. More specifically, the possibility of being last in the W&L mechanism reduces the incentive to choose a low effort, and at the same time the existence of a top prize gives rise to some higher effort choices.

Not surprisingly, the same basic pattern can be observed in contests of size three. More specifically, the distribution of effort in the winner only treatment appears bimodal in panel (a), while that in the loser and in the winner&loser treatments appears closer to a normal distribution (if somewhat right-skewed). In line with Result 3, this is a further indication that the underlying mechanism driving effort is the same in contests of size three and of size six. Notice also that effort levels in panel (a) are shifted upwards compared to panel (b); this is also consistent with the level effect mentioned in Result 3 and it is most evident in the loser tournament.

5.2 Efficiency

Hypothesis 2 states that the efficiency of each mechanism should be the same, since effort in all 6 treatments is predicted to be identical and the amount of prize money per subject is the same in all treatments. We will first focus on overall efficiency. Table 4 shows overall efficiency by treatment, defined as the combined gains to the principal and the agents.[16] The p-values in the table are from Mann-Whitney tests comparing the gains from representative groups from each treatment.[17]

As seen in Table 4, the least efficient mechanism regardless of the size of the tournament is the winner tournament, while there is no statistical difference between the winner&loser and the loser tournaments. What is also seen in this table is that actual gains per group are always below expected gains. This is due to the fact that effort is below the predicted equilibrium level, which was shown to be socially efficient in Section 3.2.

To calculate which mechanism is the most efficient from the agent's perspective, we merely need to consider the prize gained and subtract the cost to the agent of attaining the prize. Table 5 shows the expected and actual gains to the agents in each of the six treatments. The p-values in the table are from Mann-Whitney tests comparing the gains from each agent in a given treatment. As evidenced in this table, the most profitable mechanism for the agents depends on the size of the tournament. In tournaments of size six, the loser tournament is the most profitable option by quite a margin, while in tournaments of size three there is no clear ranking, although the winner tournament again appears to be the least profitable mechanism. The superiority of the LOS mechanism in tournaments of size six is due to the convex cost structure and the fact that the variance is lowest under this mechanism. It is somewhat surprising that there is no statistically significant difference between WIN and the other two mechanisms when $n = 3$ (given what appears to be a large difference), and the culprit for this is the large variance in the winner treatment, seen previously in Figure 2.

[16] As a reminder, total profit to the principal in the winner tournament is $\bar{\Pi} = n\bar{e} - V_1 - (n-1)V_2$ and payoffs for all agents in a winner tournament is $\bar{\pi} = V_1 + (n-1)V_2 - nc(\bar{e})$. This leads to the total efficiency of $n\bar{e} - nc(\bar{e})$. This calculation also holds true for LOS and for W&L tournaments.

[17] Due to random rematching of subjects into a different group each period, a representative group as used in our analysis is defined as follows: in any given period, n sujects are randomly drawn without replacement from the population of k subjects to form $\frac{k}{n}$ groups. Since there are 20 periods, a single group is the result of 20 such draws. This exercise was done to arrive at twelve groups in the $n = 3$ treatments and six groups in the $n = 6$ treatments reported in Table 4. Other methods we tried confirm the robustness of the findings in the table, but were less conservative. As an example, if the efficiency of a single group is analyzed each period in the same manner as in Table 4, the sample size is multiplied by 20 and thus stronger results are obtained. Additionally, the results remain unchanged if regressions are used where group level controls are used.

Mechanism	n	Expected Gains per group	Actual Gains per group	$p-value$ WIN vs. LOS	$p-value$ W&L vs. WIN	$p-value$ W&L vs. LOS
Winner	3	328,920	184,874			
Loser	3	328,920	260,038			
Winner&Loser	3	328,920	270,982	< 0.01	< 0.01	0.64
Winner	6	657,840	361,569			
Loser	6	657,840	505,234			
Winner&Loser	6	657,840	485,327	< 0.01	< 0.01	0.26

Table 4: Expected and actual efficiency for each treatment. The gains per group are given where a group consists of one principal and n agents. Thus, in the n=3 treatments, there were twelve groups in each treatment while in the n=6 treatments, there were six groups in each treatment. The p-values are based on Mann-Whitney tests.

Mechanism	n	Expected Gains per agent	Actual Gains per agent	$p-value$ WIN vs. LOS	$p-value$ W&L vs. WIN	$p-value$ W&L vs. LOS
Winner	3	109,640	91,722			
Loser	3	109,640	101,600			
Winner&Loser	3	109,640	100,901	0.12	0.18	0.67
Winner	6	109,640	103,221			
Loser	6	109,640	127,883			
Winner&Loser	6	109,640	103,042	< 0.01	0.81	< 0.01

Table 5: Expected and actual gains to the agents for each treatment. There were 36 observations per treatment. P-values are the result of Mann-Whitney tests.

Result 4: The winner tournament is the least efficient for tournaments of size three and size six, while there is no difference in terms of aggregate efficiency between LOS and W&L. The loser tournament is the most profitable mechanism for the agents in tournaments of size six, while there is no clear difference in tournaments of size three.

From the principal's perspective, the preferred mechanism is of course the one that generates the highest total effort. We have already dealt with this issue in Section 5.1, where we saw that the mechanism that combines a separate winner and loser prize is superior to the other two with respect to total agent effort, especially in contests of size six.

5.3 Analysis of dynamics

Results 1-4 reveal that the mechanisms considered produce quite different results. This is counter to our predictions based on the basic model of Section 3. The next step is to try and find a suitable explanation, which fits our data. Behavior observed in mechanisms such as ours has been traditionally modeled using reinforcement learning and/or directional learning. The most fundamental property of reinforcement learning is that strategies which have led to bad outcomes are less likely to be used in the future (Roth and Erev 1995). The key feature of the directional learning theory (Selten and Stoecker 1985) is that subjects respond in the direction of higher profits. If directional learning is applicable in our setting, then we should observe the following pattern: after receiving negative reinforcement (i.e., after being last in the loser tournament or after not winning a winner tournament), a subject should not choose an effort equal to or less than what led to that outcome, meaning that a higher effort will be chosen. Likewise, if an agent wins the contest or is not last, this conveys the signal that a lower number could have potentially led to higher profits; thus, a lower number will be chosen. As suggested by Grosskopf (2003), we use a combination of reinforcement and directional learning in a more formal analysis in the Appendix. Below we outline the general conjectures of reinforcement and directional learning applied to our setting.

(RD1) Following being last or not being first, subjects will increase their effort.

(RD2) Following being first or not being last, subjects will decrease their effort.

(RD3) These effects will be lessened over time.

Focusing first on the treatments that involve loser prizes (treatments LOS and W&L), a first graphical test of conjectures RD1 and RD2 is given in panel (a) of Figure 3, which examines behavior before and after being last in contests of size three and size six. Panel (a) shows that, consistent with RD1, subjects increase their effort after being last. The

differences are statistically significant both in the LOS and in the W&L treatment, and in tournaments of size three as well as size six ($p < 0.01$).[18]

Panel (b) in Figure 3 examines behavior before and after subjects learn that they were not last. We see here that, upon learning that they were not last in the previous round, subjects decrease their effort. Again, this result is statistically significant for all cases, regardless of treatment and tournament size ($p < 0.01$).

In order to check for the robustness of these findings, we present in columns (1) and (2) of Table 6 the results from two regressions, where the dependent variable is the difference in chosen effort from round $t-1$ to round t. The main explanatory variable is $LagLast$, which is a dummy variable equal to 1 if the subject was last in round $t-1$. Additionally, there is a variable accounting for the number of times a subject had already been last at the start of a given round, which is meant to capture the decreasing effect of reinforcement over time (see RD3), as well as a control variable for the size of the contest. Column (1) examines the loser contests and column (2) the winner&loser contests. For succinctness, we relegate to the Appendix specifications that use the pooled sample of LOS and W&L observations, but we note here that all our results are robust to such a specification.

In line with RD1, we see that the effect of being last in round $t-1$ causes subjects to increase their chosen effort by a sizeable amount, compared to a subject who was not last (see the large and strongly significant coefficient of $LagLast$ in both specifications). In (2) we see that, in W&L, this effect is inversely related to the number of times a subject was last, as RD3 predicts. Interestingly, however, the effectiveness of the loser tournament is not diminished in the LOS treatment.[19] Thus, these results are largely in line with learning theory.

Result 5: In contests involving single loser prizes, the evolution of play is consistent with basic learning predictions where subjects increase their effort after losing and decrease their effort following not losing.

Turning now to the effect of a single winner prize, panel (c) of Figure 3 looks at the response of agents upon winning the contest, and panel (d) examines the response upon not winning the contest. It is evident in panel (c) that, following a win, a subject decreases their effort as predicted by RD2. The decrease after winning the contest is statistically significant in the WIN as well as in the W&L contest ($p < 0.01$). Panel (d) of Figure 3

[18] Unless otherwise mentioned, all the p-values reported here are a result of a pairwise regression where errors are clustered at the subject level.

[19] This result is intriguing and the explanation for it is not given by learning theory. One explanation could hinge on what a subject focuses on. In the LOS treatment, they focus solely on whether they were last or not, while in the W&L treatment, they are also competing for the top prize. Thus, loser prizes may be more effective over time if it is the sole mechanism.

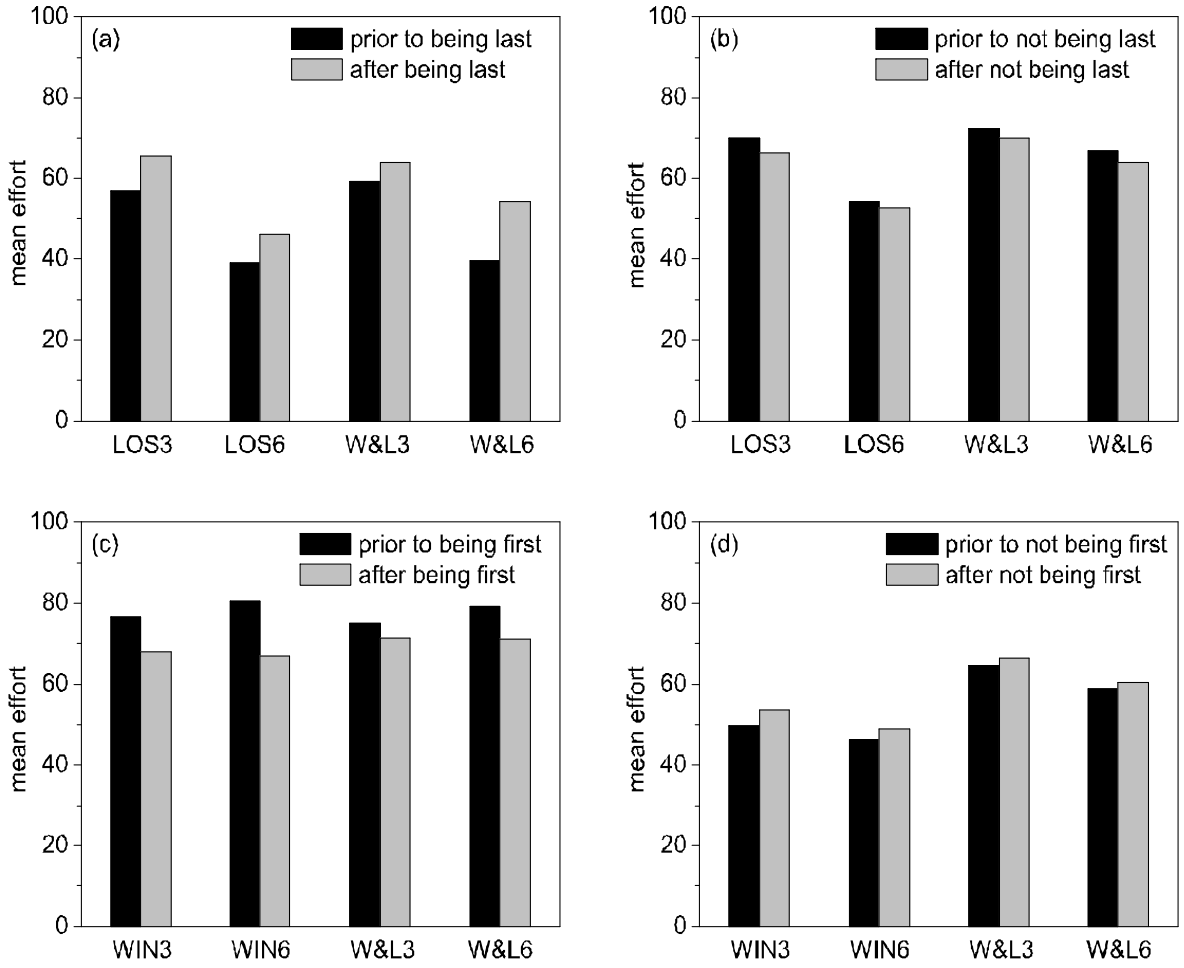

Figure 3: The reaction of subjects in tournaments before and after being reinforced. Panel (a) shows the reaction of subjects before and after being last, Panel (b) shows the reaction of subjects before and after *not* being last, Panel (c) shows the reaction of subjects before and after winning while Panel (d) shows the reaction of subjects before and after *not* winning.

Difference in Effort	**(1)** LOS	**(2)** W&L	**(3)** WIN	**(4)** W&L
Constant	−2.33	−9.36**	11.10**	2.33
	(2.74)	(3.91)	(3.24)	(3.19)
LagLast	12.38***	14.48***		
	(2.51)	(2.47)		
LagFirst			−17.30***	−10.78***
			(2.96)	(2.02)
Size Six	0.72	1.60**	−1.00	−0.91**
	(0.52)	(0.50)	(0.77)	(0.38)
Number of Times Last	−0.48	−0.90**		
	(0.59)	(0.39)		
Number of Times First			1.14**	1.07***
			(0.51)	(0.33)
Age	0.04	−0.07	−0.12	−0.13*
	(0.07)	(0.07)	(0.11)	(0.07)
Risk Aversion	−0.05	0.28**	.05	0.15
	(0.10)	(0.12)	(0.18)	(0.10)
Loss Aversion	−0.04	0.38**	−0.39*	0.24*
	(0.10)	(0.17)	(0.24)	(0.13)
Female	−0.15	−0.08	0.17	0.14
	(0.34)	(0.49)	(0.70)	(0.41)
Round	0.00	0.02	−0.19**	−0.14***
	(0.05)	(0.06)	(0.09)	(0.05)
# of observations	1368	1368	1368	1368
# of clusters	72	72	72	72
R-squared	0.09	0.06	0.03	0.03

Table 6: Individual random effects panel regressions on the difference in effort from period $t-1$ to period t where the main explanatory variable is whether the subject was last or first in period $t-1$. Because of this, period 1 is not included in the analysis. Robust standard errors, which are clustered at the individual level, are in parentheses. Three, two and one stars represent significance at the one, five and ten percent level respectively.

shows that subjects increase their effort upon not winning the contest. This is consistent with RD1. Again, the difference is highly significant for both tournament mechanisms (WIN and W&L), although the magnitude is somewhat smaller than what we saw with loser mechanisms.

Columns (3) and (4) in Table 6 examine how subjects change their behavior in contests involving a single winner prize. The dependent variable is once again the difference in chosen effort from round $t-1$ to round t. The main explanatory variable is the dummy $LagFirst$, which is equal to one if the subject was first in round $t-1$ and zero otherwise. Column (3) examines only the WIN contest and column (4) only the W&L contest.[20] Consistent with RD2 and Figure 6, we document a strong and significant negative impact of winning the contest on effort in the subsequent round. Moreover, this effect is declining with each subsequent reinforcement, as RD3 suggests.

Result 6: In contests involving single winner prizes, the evolution of play is consistent with the basic learning predictions where subjects decrease their effort in subsequent rounds after they have won and increase their effort if they did not win.

Taken together, the six results so far give us an indication of why winner&loser contests are generally more effective than the winner-only or loser-only contests. The W&L mechanism combines the reinforcing features of both contests. To explore this further, we use data from the W&L contest in order to disentangle if losing or winning has a stronger reinforcing effect (since the W&L treatment is the only one where both winning and losing were possible). Table 7 does this in a regression that includes lags of winning and losing, along with the usual control variables.[21]

As it turns out, the effect of being last has a larger impact on a subject's subsequent chosen effort, compared to the effect of being first (Wald test; $p < 0.01$). As a reminder, the variables $LagFirst$ and $LagLast$ are looking at the relative increase or decrease in effort from a subject who was previously first or last. Not only is avoiding being last statistically more effective than trying to be first, but the coefficient on $LagLast$ is more than three times greater in magnitude than the coefficient on $LagFirst$. The relatively more effective motivation of being last indicates why the outcome in the LOS3 approaches the outcome in the W&L3 contest.

[20] Once again, all results remain unchanged if we run pooled regressions that include data from treatments WIN and W&L (see Appendix).

[21] Because we want to clearly isolate the effect of being first or being last, we do not include here the number of times that a subject has been won or lost because these are inherently related to the base effect.

Difference in Effort	**(1)**
Constant	−6.20
	(3.16)
LagFirst	−3.06***
	(1.11)
LagLast	9.80***
	(1.78)
Size Six	1.21**
	(0.56)
Age	−0.10
	(0.09)
Risk Aversion	0.28**
	(0.40)
Loss Aversion	0.36**
	(0.19)
Female	−0.24
	(0.55)
Period	−0.06
	(0.05)
# of Observations	1368
# of Clusters	72
R-squared	0.06

Table 7: Individual random effects panel regressions on the difference in effort from period $t-1$ to period t in the treatment with both winners and losers where the main explanatory variables are whether the subject won or lost in period $t-1$. Because of this, period one is not included in the analysis. Robust standard errors, which are clustered at the individual level, are in parentheses. Three stars represent significance at the one percent level.

6 Discussion and Conclusion

In organizations, managers employ incentive schemes which encourage their employees to compete to be better than their cohorts and/or to avoid being the worst among them. The economic literature up to now has mainly focused on understanding how workers compete for the top prize(s). The goal of our paper has been to complement the literature by testing a baseline principal-agent tournament model that allows us to compare how agents compete for the top or avoid being last. In a laboratory experiment, we have implemented three main incentive mechanisms: winner tournament, loser tournament, and winner&loser tournament. We have also varied the size of the tournament. Although the baseline model predicts that employee effort should be the same in all treatments, our empirical results have indicated that this is not the case.

In general, no mechanism generates higher effort levels from the agents than the one which combines winner and loser prizes. We have also found that a loser tournament produces similar results to the combined mechanism in tournaments of a relatively small size (three participants), while the winner tournament and loser tournament are equivalent in terms of effort in tournaments of a relatively large size (six participants). There are two main drivers of these results. First, we have found that behavior in all mechanisms is consistent with reinforcement and direction learning (Roth and Erev 1995, Selten and Stoecker 1985). Those subjects who had previously been last (or not first) increase their effort in subsequent rounds, while those subjects who were first (or not last) decrease their effort. Second, we have found that competition away from the bottom is more effective than competition for the top at increasing subsequent effort. Taken together these two findings suggest that, as the size of the tournament decreases, the outcome in smaller contests which only incorporate a single loser prize will begin to look like contests which incorporate a winner and a loser prize.

Turning to efficiency, we have found that the winner tournament is the least efficient one from an aggregate point of view, while the efficiency achieved in the two mechanisms that include a loser prize are very similar. From the agent's perspective, the loser tournament is the most efficient (i.e., profitable) mechanism in large contests while there is no clear ordering of efficiency in the smaller contests. Since the total amount of prize money to be awarded per agent is the same in all treatments, the calculation of the most profitable mechanism for the agents depends on the total cost of expended effort. The total cost of effort is the lowest in the loser only treatments, especially in contests of size six. This is due to the much lower variance in effort in this treatment, compared to the other two. Thus, somewhat surprisingly, agents should prefer the loser tournaments.

Our paper can inform managers and policy makers on several key issues. Negative

and positive incentives are often used to reinforce good or bad behavior, and our findings indicate that they achieve this goal in a tournament setting. In order to elicit the highest amount of effort from their employees, managers should continue to use both mechanisms. Additionally, the reinforcing effect works such that the temporary use of these mechanisms will only lead to the desired outcome in the short term. Our results (see conjecture RD3) indicate that continuous reinforcement is needed in order to maintain high effort.[22] Finally, our results can better inform the actions of principals who are concerned about the high variance in performance obtained from tournaments. There are several reasons beyond efficiency why a high variance may be troubling to a principal. First, if workers exerting high effort were able to more easily observe the low effort of their co-worker(s),their best response may be to lower their effort in the long term which would result in a drastic decline in a winner tournament. Second, if having predictable quality is desireable, as is the case in many industries, a high variance in effort which results in a high variance in quality (assuming effort and quality are positively related) would be very problematic. In these settings, our results suggest principals should strongly consider using the loser tournament, as this mechanism generates the lowest variance in effort.

7 References

Akerlof, R., R. Holden. 2012. The nature of tournaments. *Economic Theory* 51(2): 289 - 313.

Balafoutas, L., E. G. Dutcher, F. Lindner, D. Ryvkin. 2012. The optimal allocation of prizes in tournaments of heterogeneous agents. University of Innsbruck working paper no. 2012-08

Bull, C., A. Schotter, K. Weigelt. 1987. Tournaments and piece rates: An experimental study. *Journal of Political Economy* 95(1): 1-33.

Camerer, C., D. Lovallo. 1999. Overconfidence and excess entry: An experimental approach. *American Economic Review* 89(1): 306-318.

Carpenter, J., P. H. Matthews, J. Schirm. 2010. Tournaments and office politics: Evidence from a real effort experiment. *American Economic Review* 100(1): 1-18.

Casas-Arce, P., F. A. Martínez-Jerez. 2009. Relative performance compensation, contests, and dynamic incentives. *Management Science* 55(8): 1306-1320.

[22] A similar result is found by Mulligan and Schaffer (2011) using simulations.

Cason, T. N., W. A. Masters, R. M. Sheremeta. 2010. Entry into winner-take-all and proportional-prize contests: An experimental study. *Journal of Public Economics* 94(9-10): 604-611.

Dechenaux, E., D. Kovenock, R. Sheremeta. 2012. A Survey of experimental research on contests, All-Pay Auctions and Tournaments. SSRN working paper

Dickinson, D. 2001. The carrot vs. the stick in work team motivation. *Experimental Economics* 4 (1): 107-124.

Eriksson, T., S. Teyssier, M. C. Villeval. 2009. Self-selection and the efficiency of tournaments. *Economic Inquiry* 47(3): 530-548.

Falk, A., E. Fehr, D. Huffman. 2008. The power and limits of tournament incentives. Discussion Paper University of Zurich.

Fischbacher, U. 2007. z-Tree: Zurich toolbox for readymade economic experiments. *Experimental Economics* 10(2): 171-178.

Gächter, S., E. Johnson, A. Herrmann. 2010. Individual-level loss aversion in riskless and risky choices. University of Nottingham Centre for Decision Research and Experimental Economics Discussion Paper, No. 2010-20.

Gneezy, U., M. Niederle, A. Rustichini. 2003. Performance in competitive environments: Gender differences. *Quarterly Journal of Economics* 118(3): 1049-1074.

Gradstein, M., K. Konrad. 1999. Orchestrating rent-seeking contests. *Economic Journal* 109: 536–545.

Greiner, B. 2004. An online recruitment system for economic experiments. In: *Forschung und wissenschaftliches Rechnen*, ed. Kremer and Macho, 79-93. Göttingen: GWD

Grosskopf, B. 2003. Reinforcement and directional learning in the ultimatum game with responder competition. *Experimental Economics* 6(2): 141-158.

Gürtler, O., M. 2011. Kräkel. Dismissal tournaments. University of Cologne working paper

Harbring, C., B. Irlenbusch. 2003. An experimental study on tournament design. *Labour Economics* 10(4): 443-464.

Harbring, C., B. Irlenbusch. 2008. How many winners are good to have? On tournaments with sabotage. *Journal of Economic Behavior and Organization* 65(3): 682-702

Harbring, C., B. Irlenbusch. 2012. Sabotage in tournaments: Evidence from a laboratory experiment. *Management Science.* 57(4): 611-627

Höchtl, W., R. Kerschbamer, R. Stracke, U. Sunde. 2011. Incentives vs. selection in promotion tournaments: Can a designer kill two birds with one stone? IZA Discussion Paper 5755.

Holt, C. A., S. K. Laury. 2002. Risk aversion and incentive effects. *American Economic Review* 92(5): 1644-1655.

Konrad, K. A. 2009. *Strategy and dynamics in contests*, Oxford University Press.

Kräkel, M. 2000. Relative deprivation in rank-order tournaments. *Labour Economics* 7(4): 385–407.

Lazear, E. P. 1995. *Personnel economics*. Cambridge, MA: MIT Press.

Lazear, E. P., S. Rosen. 1981. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* 89(5): 841-864.

List, J.A., van Soest, D., Stoop, J., & Zhou. H. 2010. On the role of group size in tournaments: Theory and evidence from lab and field experiments. Working paper, University of Chicago, Chicago.

Masclet, D., C. Noussair, S. Tucker, M. C. Villeval. 2003. Monetary and non-monetary punishment in the voluntary contributions mechanism. *American Economic Review* 93(1): 366-380.

McLaughlin, K. J. 1988. Aspects of tournament models: A survey. *Research in Labor Economics* 9: 225-256.

Mirrlees J. 1975. The Theory of moral hazard and unobservable behavior. Mimeo, Nuffield College.

Moldovanu, B., A. Sela, X. Shi. 2012. Carrots and sticks: prizes and punishments in contests. *Economic Inquiry*, 50(2): 453-462

Mulligan, J., R. B. Schaefer. 2011. A new hope for rank and yank. *Journal of Leadership and Organizational Studies* 18(3): 385–396.

Nalebuff, B. J., J. E. Stiglitz. 1983. Prizes and incentives: Towards a general theory of compensation and competition. *Bell Journal of Economics* 14(1): 21-43.

Noussair, C., S. Tucker. 2005. Combining monetary and social sanctions to promote cooperation. *Economic Inquiry* 43(3): 649-660.

O'Flaherty, B., A. Siow. 1995. Up-or-out rules in the market for lawyers. *Journal of Labor Economics* 13: 709–735.

Orrison, A., A. Schotter, K. Weigelt. 2004. Multiperson tournaments: An experimental examination. *Management Science* 50(2): 268-279.

Rosen, S. 1986. Prizes and incentives in elimination tournaments. *American Economic Review* 76(4): 701–715.

Roth, A., I. Erev. 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8(1): 164–212.

Ryvkin, D., A. Ortmann. 2008. The predictive power of three prominent tournament formats. *Management Science* 54(3): 492–504.

Schotter, A., K. Weigelt. 1992. Asymmetric tournaments, equal opportunity laws, and affirmative action: Some experimental results. *Quarterly Journal of Economics* 107(2): 511-539.

Selten, R. and Stoecker, R. 1986. End behaviour in sequences of finite prisoner's dilemma supergames: A learning theory approach. *Journal of Economic Behavior and Organization* 7(1): 47-70.

Sheremeta, R. 2010. Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior* 68(2): 731-747.

Sunde, U. 2009. Heterogeneity and performance in tournaments: a test for incentive effects using professional tennis data. *Applied Economics* 41(25): 3199-3208.

Sutter, M., S. Haigner, M. Kocher. 2010. Choosing the stick or the carrot? Endogenous institutional choice in social dilemma situations. *Review of Economic Studies* 77(4): 1540-1566.

A Experimental instructions

Instructions for WIN6:

Welcome to an experiment on decision making. We thank you for your participation!

The experiment will be conducted on the computer. All decisions and answers will remain confidential and anonymous. Please do not talk to each other during the experiment. If you have any questions, please raise your hand and we will come by and answer it.

During the experiment, you and the other participants will be asked to make a series of decisions. Your payment will be determined by your decisions as well as the decisions of the other participants according to the following rules.

During the experiment you will be earning tokens. At the end of the experiment, tokens will converted to Euros at a rate of 2000 tokens = 1 Euro. Today's experiment consists of several parts. The instructions for the first part are given below.

Rounds and Groups:

The first part consists of 20 rounds. The computer will choose 4 rounds at random for which to pay you. You will not be told which rounds will be paid until the conclusion of all parts of the experiment.

At the beginning of each round you will be randomly matched in a Group with 5 other participants. This means that in each round the groups are re-matched, so that they will not be the same (unless by chance). You will never be told the identity of those in your Group and they will never be told your identity.

Tasks:

Your task in today's experiment is to choose a number between 1 and 100. You will enter your chosen number in the blank box on your computer screen labeled "Number Chosen" and then hit "Continue." The sheet labeled "Decision Costs" shows you the cost in tokens associated with each number. Notice that higher the number chosen, the higher the associated cost. Each member in your Group has the same cost sheet as you. In each round, all Group members choose his/her numbers simultaneously. You will not know the number chosen by any of your Group members when you make your choice and likewise, they will not know the number you chose when they make their choice.

After all group members have made their choice, the computer will draw a random number between -44 and 44, independently for each member of your group. All numbers in this range are equally likely and each number drawn does not affect the number drawn for someone else in your Group. This number will be added (or subtracted) from your chosen number to make your total number.

Payoffs:

The computer will compare your total number with the total number of those in your Group. The person with the highest total number will receive 14,667 tokens while the remaining 5 members of the group will receive 5867 tokens. The cost of each chosen number will be subtracted from this amount to give you the total payment for each round should that round be chosen for payment.

At the end of each round you will be shown the random number chosen for you, your resulting total number, and whether your total number is higher than anyone's in your Group.

Example:

Let's go through an example. Suppose you chose the number 50 and the other members of your group chose 32, 65, 80, 46 and 18. Also suppose that the random number drawn for you was 26 and the random number drawn for the other members of your Group were -12, 41, -32, 13 and 7 respectively. This would mean your total number is 50+26=76. The total numbers of the other group members would be 20, 106, 48, 59 and 25. In this example, you have the second highest number and thus would receive 5900 - 633 = 5267 tokens if this round were randomly chosen for payment. Notice that the 633 tokens corresponds to the cost associated with a chosen number of 50.

If on the other hand, you had chosen 85 and all other chosen numbers and random draws remained the same, you would have a total number of 85+26=111. This would mean you would have the highest total number and would receive 14700-3789=10911 tokens if this round were randomly chosen for payment.

As a final point, once you have made your decisions or are finished viewing the results please hit the continue button. No one can move to the next round until everyone in the experiment has clicked on this button so make sure to pay attention to the screen to keep the experiment moving along.

Are there any questions?

Decision Costs:

Chosen Number	Token Cost
1	6
2	12
3	18
4	24
5	31
6	38
7	45
8	52
9	59
10	66
11	74
12	82
13	90
14	98
15	107
16	116
17	125
18	134
19	143
20	153
21	163
22	173
23	183
24	194
25	205
26	217
27	229
28	241

Chosen Number	Token Cost
29	253
30	266
31	280
32	294
33	308
34	323
35	338
36	354
37	370
38	387
39	404
40	422
41	441
42	461
43	481
44	502
45	524
46	547
47	571
48	595
49	621
50	648
51	676
52	705
53	736
54	768
55	802
56	837

Chosen Number	Token Cost
57	874
58	912
59	953
60	996
61	1041
62	1089
63	1139
64	1192
65	1248
66	1308
67	1372
68	1439
69	1511
70	1587
71	1669
72	1757
73	1851
74	1953
75	2062
76	2180
77	2307
78	2446
79	2597
80	2763
81	2944
82	3144
83	3364
84	3609

Chosen Number	Token Cost
85	3882
86	4188
87	4533
88	4925
89	5373
90	5889
91	6488
92	7192
93	8027
94	9032
95	10260
96	11789
97	13734
98	16276
99	19715
100	24571

B Learning dynamics

In this section, we describe the reinforcement and directional learning mechanisms in detail. Suppose, more formally, that in time period $t = 1$, subject i has some initial

propensity, $q_j^i(1)$, to choose a number j which results in a profit equal to $\pi_j^a(1)$.[23] The subject chooses number j if they believe that the expected profit from doing so is greater than the profit from choosing any other number, or $\pi_j^e(1) > \pi_{-j}^e(1)$. We will normalize expected profit in period t, $\pi_j^e(t)$, to be zero and let $\pi_j^r(t) = \pi_j^a(t) - \pi_j^e(t)$ be the relative profit which is more than or less than the expected profit. Thus, if actual profit is less than the subject expected ($\pi_j^r(t) < 0$), the outcome will be seen as negative reinforcement. In our setting, we will assume subjects update the propensities such that in period $t+1$, $q_j^i(t+1) = q_j^i(t) + \pi_j^r(t)$ while $q_{-j}^i(t+1) = q_{-j}^i(t)$. Thus, the probability, $p_j^i(t)$, of subject i choosing a certain number j in round t is given by the following formula.

$$p_j^i(t) = \frac{q_j^i(t)}{\sum_{h=1}^{100} q_h^i(t)} \tag{16}$$

There are two main implications to this simple model. The first is that strategies which lead to payoffs which are lower than expected will have a lower probability of being played in the future. The second implication can be seen by noticing that because of the summation term in the denominator, reinforcement has a diminishing effect over time.

The above model does a nice job of explaining the dynamics that could lead to the behavior we observe in the loser treatments when a subject is ranked last. It dictates that a subject will update their probabilities based solely on the wrong expectations of profit from choosing number j, i.e. the information they receive informs them that a *higher* number could have potentially led to a higher profit. More specifically, we can think that subjects are updating expectations such that $\pi_j^e(t+1) = \pi_j^a(t)$. What is missing from this analysis is the behavior when a subject realizes they are not last (or are first). Figure 2 hints that the reaction in the two scenarios is not symmetric. This can easily be inferred by noticing that if the reactions to the two scenarios were symmetric, the average would approach zero since there are more people not last than are last. The direction of updating, however, is different in this scenario. If a subject realizes they are not last, this does not mean they had wrong expectations about the payoff they would receive from choosing number j; it means they had the wrong expectations about the payoffs from the other strategies, i.e., the information they receive informs them that a *lower* number could have potentially led to a higher profit. More formally, we can say that a subject's expected payoff from choosing any number is a function of their beliefs about the payoffs they can receive from any other number j. In other words, $\pi_{-j}^e(t) = \alpha(t)\pi_j^e(t)$ where α

[23] We will assume the initial propensity is fixed and will not explore what may cause a subject to develop initial propensities.

is a parameter accounting for the distance between j and other numbers close to j. The closer the number is to j, α approaches 1 from above or below. Notice that if a player thinks that choosing number j is optimal, then $\alpha < 1$.[24] In each period then, a subject updates their expected profit from choosing any number and they update how this profit relates to similarly chosen efforts.

The implications of a structure where there exists a correlation in profits as defined is obvious and results in the asymmetries observed. This can be explained by noticing that by choosing a higher number in period t after losing in period $t-1$ results in a higher cost, but these costs are offset by the much larger gain attainable if the subject is not last. This is in contrast to a subject who finds they are not last. The gain from choosing a lower number is small since it is only a cost savings, but the potential loss if the subject is last is quite large. Or, put more simply, the potential payoff gain in period t from large deviations from the number chosen in period $t-1$ is much greater if the subject was last in period $t-1$ than if they were not last.

C Robustness checks

Below are regressions run with the full samples which serve as robustness checks for Table 6.

[24] This is similar in spirit to the "experimentation" parameter in the Roth-Erev Model.

Difference in Effort	**(1)** LOS	**(2)** W&L	**(3)** Pooled	**(4)** Pooled
Constant	−2.33	−9.36**	−5.37**	−5.07**
	(2.74)	(3.91)	(2.25)	(2.35)
LagLast	12.38***	14.48***	13.52***	11.94***
	(2.51)	(2.47)	(1.72)	(2.60)
W&L			−0.01	−0.07
			(.29)	(0.66)
Size Six	0.72	1.60**	0.99***	0.42
	(0.52)	(0.50)	(0.34)	(0.66)
Number of Times Last	−0.48	−0.90**	−0.74**	−0.63*
	(0.59)	(0.39)	(0.32)	(0.34)
Lag Last×Size Six				2.65
				(2.50)
Lag Last×W&L				0.28
				(2.30)
Age	0.04	−0.07	0.01	0.01
	(0.07)	(0.07)	(0.05)	(0.05)
Risk Aversion	−0.05	0.28**	0.08	0.07
	(0.10)	(0.12)	(0.07)	(0.08)
Loss Aversion	−0.04	0.38**	0.10	0.12
	(0.10)	(0.17)	(0.09)	(0.10)
Female	−0.15	−0.08	0.00	0.06
	(0.34)	(0.49)	(0.29)	(0.31)
Period	0.00	0.02	0.02	0.00
	(0.05)	(0.06)	(0.04)	(0.04)
# of Observations	1368	1368	2736	2736
# of Clusters	72	72	144	144
r-squared	0.09	0.06	0.07	0.07

Table 8: Random effects panel regressions on the difference in effort from period $t-1$ to period t where the main explanatory variable is whether the subject was last in period $t-1$. Because of this, period 1 is not included in the analysis. Robust standard errors, which are clustered at the individual level, are in parenthesis. Three, two and one stars represent significance at the one, five and ten percent level respectively. Columns (1) and (2) are given in the text while columns (3) and (4) are used as robustness checks.

Difference in Effort	**(1)** WIN	**(2)** W&L	**(3)** Pooled	**(4)** Pooled
Constant	11.10**	2.33	6.38**	6.40**
	(3.24)	(3.19)	(3.09)	(3.28)
Lag First	−17.30***	−10.78***	−13.82***	−16.00***
	(2.96)	(2.02)	(1.787)	(2.71)
W&L			0.49	−1.06
			(0.36)	(0.71)
Size Six	−1.00	−0.91**	−1.10***	−0.66
	(0.77)	(0.38)	(0.42)	(0.73)
Number of Times First	1.14**	1.07***	1.04***	1.00***
	(0.51)	(0.33)	(0.31)	(0.32)
Lag First×Size Six				−2.02
				(2.72)
Lag First×W&L				6.19***
				(2.40)
Age	−0.12	−0.13*	−0.13**	−0.11
	(0.11)	(0.07)	(0.07)	(0.07)
Risk Aversion	.05	0.15	0.10	0.11
	(0.18)	(0.10)	(0.09)	(0.10)
Loss Aversion	−0.39*	0.24*	−0.03	−0.05
	(0.24)	(0.13)	(0.12)	(0.13)
Female	0.17	0.14	0.02	0.08
	(0.70)	(0.41)	(0.39)	(0.41)
Period	−0.19**	−0.14***	−0.16***	−0.16***
	(0.09)	(0.05)	(0.05)	(0.05)
# of Observations	1368	1368	2736	2736
# of Clusters	72	72	144	144
r-squared	0.03	0.03	0.03	0.03

Table 9: Random effects panel regressions on the difference in effort from period $t-1$ to period t where the main explanatory variable is whether the subject was first in period $t-1$. Because of this, period 1 is not included in the analysis. Robust standard errors, which are clustered at the individual level, are in parenthesis. Three, two and one stars represent significance at the one, five and ten percent level respectively. The models in columns (1) and (2) are given in the text while the models in columns (3) and (4) are used as robustness checks.