

5-18-2016

Dynamic Behavior and Player Types in Majoritarian Multi-Battle Contests

Alan Gelder

Chapman University, geldera@chapman.edu

Dan Kovenock

Chapman University, kovenock@chapman.edu

Follow this and additional works at: http://digitalcommons.chapman.edu/esi_working_papers



Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Gelder, A., & Kovenock, D. (2016). Dynamic behavior and player types in majoritarian multi-battle contests. ESI Working Paper 16-12. Retrieved from http://digitalcommons.chapman.edu/esi_working_papers/190

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Dynamic Behavior and Player Types in Majoritarian Multi-Battle Contests

Comments

Working Paper 16-12

Dynamic Behavior and Player Types in Majoritarian Multi-Battle Contests

Alan Gelder and Dan Kovenock*

Economic Science Institute, Chapman University, Orange, CA 92866, USA

Abstract

In a dynamic contest where it is costly to compete, a player who is behind must decide whether to surrender or to keep fighting in the face of bleak odds. We experimentally examine the game theoretic prediction of last stand behavior in a multi-battle contest with a winning prize and losing penalty, as well as the contrasting prediction of surrendering in the corresponding contest with no penalty. We find varied evidence in support of these hypotheses in the aggregated data, but more conclusive evidence when scrutinizing individual player behavior. Players' realized strategies tend to conform to one of several "types". We develop a taxonomy to classify player types and study how these types interact and how their incidence varies across treatments. Contrary to the theoretical prediction, escalation is the predominant behavior, but last stand and surrendering behaviors also arise at rates responsive to the importance of losing penalties.

Keywords: Dynamic Contest, Multi-Battle Contest, Player Type, Experiment, All-Pay Auction, Escalation, Last Stand, Maximin
JEL: C73, C92, D44, D72, D74

1. Introduction

Contests are commonly dynamic in nature. Businesses may spend weeks or months jockeying for a lucrative contract, political campaigns frequently span a year or more, and major wars span several. Just as in war, a unifying feature of many contests is that they are frequently comprised of several smaller battles, and that

[☆]This is an updated version of Economic Science Institute Working Paper No. 15-02 (SSRN No. 2558229), dated 27 January 2015. An even earlier version of this paper circulated under the title *Fight or Surrender: Experimental Analysis of Last Stand Behavior*.

*Corresponding author. Phone: +1 (714) 628-7226

Email addresses: gelder@chapman.edu (Alan Gelder), kovenock@chapman.edu (Dan Kovenock)

throughout those battles, participants have an awareness of whether they are falling behind or approaching an overall victory. Such contests are also inherently costly—whether the cost is denominated in troops, effort, money, or other form of expenditure. Participants must therefore weigh the amount and timing of these expenditures in view of their proximity to victory or defeat and the associated consequences of ultimately winning or losing. Needless to say, a broad host of strategies may ensue.

Building on recent theoretical work on dynamic contests, we experimentally examine strategic behavior in the two-player best-of-seven tournament with complete information. Battles within the tournament occur sequentially and are modeled as all-pay auctions. That is, the high bidder wins, but both players pay the cost of their own bid (see Hillman and Riley 1989; and Baye et al. 1996). The first player to win four all-pay auctions becomes the tournament winner and receives a positive winning prize while the loser incurs a nonpositive losing penalty.¹ Under the assumption that the losing penalty is zero and that there is no time discounting of tournament winnings or bidding expenditures in the all-pay auctions, Konrad and Kovenock (2009) find that the unique sub-game perfect equilibrium entails intense competition when players are symmetrically situated in the tournament. But when a player falls behind by even one battle, he is expected to completely give up.² Gelder (2014) shows that this stark pattern of surrendering is not robust to the combined inclusion of a strictly negative losing penalty and the time discounting of potential future earnings or losses. Rather, a player who is behind in the tournament by two or more battles may exhibit last stand behavior—a phenomenon characterized by the laggard bidding more aggressively on average than the front-runner, despite the poor odds of an ultimate tournament win. A last stand is much more the story of delaying imminent loss than the story of gaining the victor’s crown.

¹This type of contest is a special case of what is known more generally as a two-player race, a term that stems back to the patent race model of Harris and Vickers (1987). Instead of using the all-pay auction, Harris and Vickers model the component contests with a logit-style lottery contest success function where a player’s probability of winning is the ratio of the player’s own bid to the sum of both players’ bids. Harris and Vickers additionally scale the cost of a player’s bid by the sum of the bids.

²Konrad and Kovenock’s analysis covers a more general class of multi-battle contests that allows for the number of all-pay auction victories needed to win the tournament to vary across players and for players to accrue asymmetric prizes for winning the tournament and identical prizes for individual battle victories. If prizes for individual battles are present, a player who is behind will compete, but he will essentially behave as if he were solely competing for the individual battle prize and not for the overall winning prize. An early example of a dynamic contest where players slacken their effort or give up entirely if they fall behind is Fudenberg et al.’s (1983) study of preemption in patent races by firms who have a marginal lead over their competitors.

We experimentally investigate the contrasting predictions of surrendering versus making a last stand using three prize-penalty combinations: one in which there is a large prize but no penalty, one with a prize and penalty of equal magnitude, and a third in which the penalty is dramatically larger than the prize. In each case, the net difference between the prize and penalty is identical so that in a one-shot contest with risk neutral players, equilibrium behavior would be the same in each case. Because discounting is a central ingredient for the last stand predictions of Gelder (2014), we implement discounting via a probability that the best-of-seven tournament suddenly terminates after each battle. In which case, neither player receives the winning prize or losing penalty, although previous bids are still forfeited.³

In accordance with the theoretical predictions, along the path of a best-of-seven tournament where one player is consistently losing each battle, the rate of surrender is conspicuously high in the treatment with no losing penalty. Along that same path, last stand behavior is a salient feature of the treatments with losing penalties. Viewed from the full set of possible paths, however, the most prominent behavior in all treatments is to escalate the conflict by engaging in a bidding war. The combination of aggressive escalation, passive surrendering, and an artful last stand which becomes aggressive at key moments results in tournaments that have a certain degree of signaling during the initial battles. Although the equilibrium bidding strategy at each battle is entirely forward looking—eyeing each player’s distance from a tournament win, rather than the particular bid realizations and path which lead to that point—players clearly respond to the bidding history in the experiment.

Our analysis is twofold, and begins by looking at aggregate behavior on a state-by-state basis within the best-of-seven tournament. Although the dynamic analysis at this stage is largely limited to comparisons between the previous and current battle, the patterns of escalation, surrender, and last stand behavior are still discernible. We next scrutinize dynamic behavior further by examining the full path of each tournament. Here we find that players gravitate toward different strategies and repeatedly use them throughout the experiment. In this regard, players can be categorized according to one of several player types. Since a player’s full extensive form strategy is not available in the experimental data, which only contain the player’s history-contingent actions along the realized path of the tournament, we refer to these observable actions as a player’s *realized strategy*.⁴ Defining a formal taxonomy, we classify the occurrence of different realized strategies across treatments.

³Under risk neutrality, this adaptation of the tournament is theoretically innocuous.

⁴The tournament’s complex extensive form makes the strategy method (Selten 1967) difficult to implement without strong and unrealistic restrictions on the information available to players at each stage of the game.

Our paper fits within a growing literature on contests and tournaments (Dechenaux et al. 2015 provides an extensive survey of contest experiments), as well as within the behavioral literature on heterogeneous player types in experiments. In terms of dynamic contest experiments, we bridge the gap between several best-of-three tournaments (e.g. Sheremeta 2010; Mago and Sheremeta 2012; Mago et al. 2013; and Irfanoglu et al. 2014) and the best-of-19 tournament in Zizzo (2002). Deck and Sheremeta’s (2012) game of siege is also closely related as it can be reached as an intermediate stage in a best-of-seven tournament. In their experiment, players are positioned asymmetrically so that one player (the defender) needs to win three or four successive battles to be victorious, while the attacker only needs to win one. Instead of automatically positioning one player on the brink of defeat, our experiment allows players to reach that point endogenously. It also permits us to see the history of bids leading to that point.⁵

In terms of the player types literature, subjects approach experiments with varying degrees of strategic sophistication, and *level-k* and other theories attempt to classify this range of behavior.⁶ Other experimental work induces player types by allowing subjects to potentially be matched against a computer which is known to be programmed a certain way (e.g. Embrey et al. 2015). Here, rather than inducing player types or focusing on the degree to which players are best-responding, we define a handful of player types corresponding to common endogenous patterns of play and analyze their frequency. We also provide a rough ranking of these types in terms of average payoffs.

2. Theory and Hypotheses

The winner of a best-of-seven tournament is the first player to win four battles. To track each player’s progress, we model the state space as a pair (i, j) where i is the number of battles that Player A still needs to win and j is the number of battles that Player B still needs to win.⁷ Hence, the tournament begins at state $(4, 4)$ and

⁵As we will see, the history of bids is especially revealing. Some players surrender from the start, others begin by mimicking a surrender and then make a last stand. Still others are clearly aggressive but happen to consistently lose by bad luck. Although Deck and Sheremeta avoid any endogenous self-selection by automatically starting players at this point, they miss both the history and the ability to observe competition at less crucial battles.

⁶A recent survey on level-k, cognitive hierarchy, and other related theories is Crawford et al. (2013). Fragiadakis et al. (2013) assert that while strategic play can be captured by such theories, more needs to be done in modeling non-strategic play which follows replicable rules-of-thumb.

⁷Tracking the absolute number of wins for each player requires a two dimensional state space. An alternative model, known as the tug-of-war, tracks the relative number of wins with a unidi-

	B wins					
0	(4, 0)	(3, 0)	(2, 0)	(1, 0)		
1	(4, 1)	(3, 1)	(2, 1)	(1, 1)	(0, 1)	
2	(4, 2)	(3, 2)	(2, 2)	(1, 2)	(0, 2)	A wins
3	(4, 3)	(3, 3)	(2, 3)	(1, 3)	(0, 3)	
4	(4, 4)	(3, 4)	(2, 4)	(1, 4)	(0, 4)	
	4	3	2	1	0	

Figure 1: Best-of-seven tournament

proceeds until it reaches $(0, j)$ for $(i, 0)$ for $i, j \in \{1, 2, 3, 4\}$. This is depicted in Figure 1. Once a player has won four battles, he receives a prize $Z \geq 0$ and his opponent incurs a penalty $L \leq 0$. Each battle consists of players competing in an all-pay auction with the winner of the auction advancing one state closer to victory.⁸ The unique equilibrium of the two-player all-pay auction is in mixed strategies with players randomizing their bids between 0 and the smaller of the two players' valuation of the prize (Baye et al. 1996). While both players randomize over this interval, the player with the lower valuation will bid 0 with positive probability. That is, if ζ_H and ζ_L are the high and low valuations of the prize ($\zeta_H \geq \zeta_L > 0$), then the equilibrium bidding distributions are as follows:

$$\begin{aligned}
 F_H(h) &= \begin{cases} h/\zeta_L & \text{if } h \in [0, \zeta_L] \\ 1 & \text{if } h > \zeta_L \end{cases} \\
 G_L(\ell) &= \begin{cases} (\zeta_H - \zeta_L + \ell)/\zeta_H & \text{if } \ell \in [0, \zeta_L] \\ 1 & \text{if } \ell > \zeta_L \end{cases} \quad (1)
 \end{aligned}$$

Expected payoffs are then $u_H = \zeta_H - \zeta_L$ and $u_L = 0$; winning probabilities are $p_H = 1 - \frac{\zeta_L}{2\zeta_H}$ and $p_L = \frac{\zeta_L}{2\zeta_H}$; and the expected bids are $\mathbb{E}[e_H] = \frac{\zeta_L}{2}$ and $\mathbb{E}[e_L] = \frac{\zeta_L^2}{2\zeta_H}$.

The bulk of the analysis in Konrad and Kovenock (2009) and in Gelder (2014) is in extending the one-shot all-pay auction to a dynamic structure where an actual

dimensional state space. Within the tug-of-war, Konrad and Kovenock (2005) predict that laggards surrender when there is no losing penalty, while Agastya and McAfee (2006) find that last stand behavior is possible when there is a penalty.

⁸An arbitrary tie-breaking rule typically suffices, but the equilibrium in Konrad and Kovenock (2009) requires that ties be awarded to the player who is ahead in the tournament. This allows the frontrunner to coast to victory by bidding zero when the laggard surrenders. Since this is a rather technical requirement, we use a fair randomizing device to break ties in the experiment.

prize is awarded only after a player achieves a critical number of wins. Hence, it becomes necessary to identify the prize valuations at each non-terminal state (i, j) for $i, j > 0$. These prize valuations are implicitly defined based on the marginal benefit of winning at (i, j) and being one state closer to overall victory versus losing and being one state closer to defeat. When losing is costless—as in Konrad and Kovenock—a player who is behind has a prize valuation of zero, so there is no incentive to compete.⁹ When there is a cost to losing and when players would prefer to win early and lose late, the prize valuations are always strictly positive, so players actively compete at all non-terminal states.¹⁰ The magnitudes of the prize valuations do, however, vary from state to state and across players. Gelder finds that there is a collection of states where the player who is behind in the tournament actually has the higher prize valuation and therefore tends to compete more aggressively. This heightened competition from the underdog is what Gelder terms the last stand.

With regard to incentives, the last stand represents the position in the tournament where the underdog’s incentive to avoid losing is stronger than the frontrunner’s incentive to win. A player who must avoid losing today, or else incur a sufficiently large penalty, has a stronger motive to compete than the opposing player who may secure the victory tomorrow if not today. The precise collection of states where a last stand occurs depends on the ratio of the winning prize to the losing penalty, as well as the discount factor. The larger the penalty, the closer to the end of the tournament the last stand occurs. In addition to the last stand, Gelder also finds that the frontrunner will defend his overall lead in the tournament if it is threatened. The *defense of the lead* occurs when the frontrunner only has a one-state lead in the tournament and entails a much higher expenditure from the frontrunner than the underdog in expectation. Thus the last stand acts as a defensive push, while the defense of the lead acts as an offensive one.

From the theoretical predictions, there are five main hypotheses we will examine. The first three address the last stand, surrendering, and the defense of the lead. The fourth examines winning margins, which are intimately connected with the conflicting behaviors of making a last stand or surrendering. The final hypothesis addresses the role of the initial battle as a predictor for the final tournament outcome.

⁹A player who is behind receives zero from continuing to lose. Since that player’s expected payoff from winning a single state is also zero, then the prize valuation is zero as well.

¹⁰An example of when these assumptions may be satisfied is the US presidential primaries. Candidates would typically prefer to secure their party’s nomination early in the election cycle to have more time to prepare for the general election. On the losing side, the potential loss of political capital is likely higher for candidates who unmistakably lose at an early stage and are not able to demonstrate their viability for future campaigns.

- H1.** *Players on a losing trajectory will make a last stand if the penalty for losing is large relative to the winning prize.*
- H2.** *With no losing penalty, a player who is behind will surrender.*
- H3.** *Players with a one battle lead in the tournament will compete more aggressively than their opponent in order to defend their lead.*
- H4.** *The expected winning margin is increasing in the size of the winning prize relative to the losing penalty.*
- H5.** *The initial battle winner wins the entire tournament more than would be predicted by an equal probability of winning each remaining battle.*

3. Methodology

We conducted 18 experimental sessions, each composed of 12 subjects, at the Economic Science Institute, Chapman University. Sessions were held in a computer lab with partitions between computers. The experiment began with subjects reading the instructions on their computer and completing a short comprehension quiz (a copy is included in the appendix). This quiz was followed by a short risk preference lottery *à la* Holt and Laury (2002). The main portion of the experiment consisted of 20 best-of-seven tournaments, and subjects were randomly and anonymously paired via the computer network for each of these. At the conclusion of the experiment, subjects completed a demographics survey and were paid in cash based on their performance in two randomly selected tournaments.

Each battle within the best-of-seven tournaments was an all-pay auction: subjects placed bids simultaneously, the high bidder won, and ties were broken randomly. In all-pay fashion, a subject's bids throughout a tournament were deducted from his or her payoff for that tournament. The tournament winner then received a prize and the loser incurred a penalty. To capture time preferences for winning and losing, we followed a common practice from macroeconomic experiments by implementing discounting via a continuation probability (e.g. Duffy 2008; Noussair and Matheny 2000). Until a player successfully won four battles, there was a 90% probability that the tournament would indeed continue from one battle to the next (as a discount factor, $\delta = 0.9$). If a tournament ended prematurely, neither player would receive a prize or a penalty, but players still had to pay their bids. Under risk neutrality, such a continuation probability is equivalent to discounting in terms of expected payoffs.¹¹

¹¹The random ending rule may also be thought of as the potential that some exogenous factor

We conducted three separate payoff scenarios, with six sessions per treatment:¹² one with a substantial losing penalty and meager winning prize (Win 15 Lose 285), the second with an equal prize and penalty (Win 150 Lose 150), and the third with a sizable prize but no penalty (Win 300 Lose 0).¹³ We abbreviate these as L285, L150, and L0. The prizes, penalties, and bids were denominated in an experimental currency called rupees, at 50 rupees = 1 US dollar. In order to make the stakes comparable across treatments, we fixed the difference between the positive prize and the negative penalty at 300 rupees. During a best-of-seven tournament, subjects could see both their own and their opponent’s previous bids.¹⁴ They also could see how many rounds they had won or lost, as well as the sum of their bids up to that point in the tournament. An example of the bidding screen is shown in the appendix. When a tournament finished, either by a player winning four battles or by an early termination, the screen would display the final outcome and payoffs before randomly re-matching subjects to begin a new tournament.

For each tournament, subjects received an endowment of rupees from which bids and the losing penalty could be deducted, and to which the winning prize could be added. The final balance was the payoff for the tournament. Since losing penalties varied across treatments, and since bids and losing penalties were deducted from the same account, we wanted to make the treatments comparable in terms of the underlying bidding budget. We accomplished this by varying the initial endowment so that it was composed of an effective bidding budget (700 rupees) plus the size of the losing penalty. Thus, for penalties of 285, 150, and 0, the endowment was 985, 850, and 700. Our choice of 700 for an effective bidding budget was an attempt

suddenly disrupts the conflict (such as the cavalry coming to save the day). Alternatively, with a present value interpretation of discounting, the size of the prize and penalty could be contingent on the winning margin. A benefit of using the random ending rule is that the order of magnitude of expenditures early in the tournament remains comparable to that of the prize and penalty at later stages of the tournament. Noussair and Matheny (2000) compared both the random ending rule and the present value interpretation of discounting in an experiment involving a single agent dynamic optimization problem and found similar results with each method.

¹²For each treatment, three of the six sessions were limited to subjects who had previously participated in some other unrelated contest experiment; the other three sessions were open to all individuals in the subject pool, including those who had been involved in another contest experiment. The results are largely unaffected by this variation in the subject pool.

¹³Experimental parameters must be multiplied by $1/\delta$ to coincide with the theoretical ones. Players in the experiment receive the prize or penalty with certainty after winning or losing four battles, but the theoretical model still has discounting between $(i, 1)$ and $(i, 0)$ (or $(1, j)$ and $(0, j)$).

¹⁴Since the unique subgame perfect equilibrium is also Markov perfect, equilibrium bids are not affected by the knowledge of bid realizations at previous states. We chose not to exogenously impose Markovian behavior in the experiment—a choice that results in an extensive form game that is too complex for utilizing the strategy method pioneered by Selten (1967).

to balance two opposing constraints: having a large enough budget so that players would not feel constrained, especially if they reach the sixth or seventh battle; but small enough so that the losing penalties have some bite.¹⁵ In each battle, players could submit bids between 0 and 300 inclusive with up to one decimal place.¹⁶ Our analysis centers on the last ten of the twenty tournaments.

4. Initial Results

We begin with a brief bird’s-eye view of bidding behavior by treatment and state. Following the format in Figure 1 of portraying the different states of a best-of-seven tournament with a four-by-four grid, the upper panel of Figure 2 plots the empirical bidding distributions from the experiment, while the bottom panel plots the underlying theoretical distributions. Along the diagonal where the tournament is tied—running from the initial state of (4, 4) at the bottom-left to the final possible state of (1, 1) at the top-right—competition becomes increasingly intense in terms of expected bids. This is true of all treatments, both theoretically and in the lab; although the bidding levels are considerably lower in the lab at these states.¹⁷

Moving away from states where the tournament is tied, the *defense of the lead* (H3) is depicted theoretically by the rather timid distributions at (4, 3), (3, 2) and (2, 1), which are dwarfed by their counterparts at (3, 4), (2, 3), and (1, 2). This timidity serves to further accentuate the *last stand* (H1) at the upper-left states of (4, 2), (4, 1), and (3, 1). Despite being two or three states behind, the expected winning probability is over one-half at each of these states in the Win 15 Lose 285 treatment, and over one-quarter in the Win 150 Lose 150 treatment.¹⁸ Finally, the tendency

¹⁵Theoretical expected expenditures throughout the entire tournament are as follows: 131.3 (Win 15 Lose 285), 114.4 (Win 150 Lose 150), and 109.4 (Win 300 Lose 0). The effective bidding budget of 700 is large enough to amply cover these amounts. However, in the theoretical model, if a player happened to consistently bid at the top of the equilibrium bidding distribution along the most expensive path of the tournament, then cumulative expenditures could be as high as 975.3 (Win 15 Lose 285), 1001.7 (Win 150 Lose 150), or 1031.7 (Win 300 Lose 0). Thus there are contingencies of the tournament for which the effective bidding budget is theoretically binding.

¹⁶300 is the highest upper bound of an equilibrium bidding distribution, occurring at (1, 1).

¹⁷In previous one-shot contest experiments, bidding distributions are frequently bifurcated with subjects submitting either high or low bids. This results in distributions that are concave over the lower bidding range and convex over the upper range (e.g. Potters et al. 1998; Gneezy and Smorodinsky 2006; and Ernst and Thöni 2013). Both for all-pay auction and Tullock contest experiments, prospect theory has been proposed as a possible explanation (Ernst and Thöni 2013; Sheremeta 2013). Here, however, bidding distributions are concave at most states, with exceptions occurring when both players are close to victory.

¹⁸Specifically, at (4, 2), (4, 1), and (3, 1), the theoretical winning probabilities are 0.529, 0.564, and 0.524 in the Win 15 Lose 285 treatment, and 0.278, 0.263, and 0.263 in the Win 150 Lose 150




Figure 2: CDFs of bids by treatment at each state: empirical (top) and theoretical (bottom).

to utterly *surrender* (H2) in the Win 300 Lose 0 treatment is represented in the theoretical distributions by the lack of any positive bids once a player falls behind.¹⁹

On quick inspection, each of these hypothesized behaviors is evidently lacking to one degree or another in the empirical bidding distributions. The defense of the lead is nowhere to be seen, with winning probabilities at (4, 3), (3, 2) and (2, 1) sitting between one-half and one-third across treatments.²⁰ When a last stand is expected, players do win 33 to 45 percent of the time, but the apparent lack of a defense of the lead masks the extent to which these wins may be interpreted as a last stand. As for surrendering, the clearest example is at (4, 1), where 74% of the bids are zero in the Win 300 Lose 0 treatment (compared with 41% in each of the other two treatments).²¹ Another potential indication of surrendering is that the Win 300 Lose 0 treatment has the lowest winning probabilities at (4, 2), (4, 1), and (3, 1).²² Yet looking across states, the bidding distributions are often more aggressive in the Win 300 Lose 0 treatment than in the Win 150 Lose 150 treatment. Surrendering, along with the other hypothesized behaviors, is not as prevalent as predicted.

To understand surrendering or last stand behavior, it is worth drilling down into the observations at state (4, 1) to explore the particular bidding histories leading to this point. Table 1 divides the state (4, 1) observations into five such histories. Each division is based on whether the player bids zero at (4, 1), as well as on the frequency of zero bids at the previous three states: (4, 2), (4, 3), and (4, 4). A bid of zero at this final state ensures defeat and can definitively be interpreted as surrendering. Bidding zero at the earlier states, however, is a mixed signal. It may indeed convey an intent to surrender. But it may also imply that the player is lying in wait—hoping perhaps that the tournament will terminate prematurely after one

treatment (for the latter treatment, a best-of-seven tournament is not large enough to include the states where the player who is behind is expected to win with more than one-half probability). In contrast, theoretical winning probabilities at (4, 3), (3, 2) and (2, 1) are 0.056, 0.056, and 0.052 in the Win 15 Lose 285 treatment, and 0.029, 0.028, and 0.026 in the Win 150 Lose 150 treatment.

¹⁹As previously noted, Konrad and Kovenock (2009) award ties to the player who is ahead in the tournament. Once the laggard surrenders, the frontrunner has no need to place a positive bid.

²⁰Namely, 0.38, 0.48, and 0.48 respectively in the Win 15 Lose 285 treatment; 0.28, 0.39, and 0.48 in the Win 150 Lose 150 treatment; and 0.32, 0.44, and 0.40 in the Win 300 Lose 0 treatment.

²¹Combining bids of 0 and 0.1 raises these totals to 78% in the Win 300 Lose 0 treatment, 45% in the Win 150 Lose 150 treatment, and 43% in the Win 15 Lose 285 treatment.

²²Experimental winning probabilities at (4, 2), (4, 1), and (3, 1) are 0.26, 0.16, and 0.33 in the Win 300 Lose 0 treatment. The Win 15 Lose 285 treatment has corresponding values of 0.33, 0.34, and 0.43, while the Win 150 Lose 150 treatment comes in at 0.33, 0.39, and 0.45. Using a standard proportions test, each of these states has a lower winning probability in the Win 300 Lose 0 treatment than each of the other two treatments at the 10% significance level or below. This comparison holds at (4, 1) well below the 1% significance level.

Table 1: Bidding history of state (4, 1) observations (in %)

History	L285	L150	L0
Bid 0 at (4, 1): And bid 0 at all previous states	23.3	25.8	50.4
Bid 0 at (4, 1): One or more previous bids > 0	17.5	15.2	23.2
Bid > 0 at (4, 1): And bid > 0 at all previous states	25.8	32.6	15.2
Bid > 0 at (4, 1): All previous bids were 0	16.7	9.8	4.8
Bid > 0 at (4, 1): One or two previous bids were 0	16.7	16.7	6.4
<i>Total obs.</i>	120	132	125

of the early rounds, but ready to act if it does not.

In the first history of Table 1, players surrender from the start and submit zero bids at each and every state. Notably, such an outright surrender accounts for half the observations at (4, 1) in the Win 300 Lose 0 treatment, twice the frequency of the other two treatments.²³ The second history also exhibits surrendering, but one in which the player was actively competing at a previous point in the tournament. This type of surrendering is perhaps more indicative of discouragement, and the difference across treatments is less stark (although still highest in the Win 300 Lose 0 treatment). The next three histories all entail a strictly positive bid at (4, 1). Of these, the most prevalent is to place positive bids at all previous states. Although this distinction does not account for changes in the magnitude of past bids, this history would certainly capture those who never gave up.²⁴ The fourth and fifth histories provide some evidence of last stand behavior. These players lie in wait, placing bids of zero at one, two, or even all three previous states, yet spring into action when facing a tournament loss. Combined, these two histories account for a full third of the (4, 1) observations in the Win 15 Lose 285 treatment, over a quarter in the Win 150 Lose 150 treatment, but only 11% in the Win 300 Lose 0 treatment.²⁵

²³A proportions test between the frequency of outright surrender in the Win 300 Lose 0 treatment and each of the other two treatments yields a p-value of 0.000.

²⁴One way to address changes in magnitude is to track whether players bid at least as high as the bid they lost to in the previous round. In the Win 15 Lose 285 and the Win 150 Lose 150 treatments, 74% of the third history observations had a bid at (4, 1) that was higher than the opponent's previous bid. For the Win 300 Lose 0 treatment, it drops to 58%.

²⁵Whether the fourth and fifth histories are combined or considered separately, the proportion of observations in these histories is statistically greater in both the Win 15 Lose 285 and the Win 150 Lose 150 treatments than in the Win 300 Lose 0 treatment. The least significant comparison is between the Win 150 Lose 150 and the Win 300 Lose 0 treatments in the fourth history (z-stat=2.39, p-value=0.061). Otherwise, all comparisons are significant below the 1% level.

Table 2: Propensity to match or outbid the winning bid in the previous state (in %)

After arriving at a state by losing					
	(4,3)	(4,2)	(4,1)	(3,1)	All Other States
L285	48.4	46.9	47.5	59.7	79.2 to 87.9
L150	40.9	40.4	47.7	63.2	67.6 to 91.7
L0	45.1	32.6	20.8	51.4	66.7 to 91.7

After arriving at a state by winning					
	(3,4)	(2,4)	(1,4)	(1,3)	All Other States
L285	66.6	60.7	56.7	72.7	80.2 to 90.9
L150	67.5	56.8	50.8	64.5	68.0 to 91.7
L0	72.1	56.5	57.6	75.7	73.6 to 89.6

We gain another clue into surrendering and last stand behavior by examining the propensity to submit bids that are at least as high as the winning bid in the previous state. These are bids which up the ante. Table 2 shows the propensity for such bids, both by state and by whether the player arrived at that state by winning or losing the previous round. The response across treatments at (4, 3) to match or outbid an opponent’s winning bid at (4, 4) ranges from 41% to 48%. Subsequent propensities at (4, 2) and (4, 1) are revealing. The Win 15 Lose 285 treatment maintains a steady course at roughly 47%; the Win 150 Lose 150 treatment swells from 40% to 48% as it reaches (4, 1); the Win 300 Lose 0 treatment conversely sinks to 21%.²⁶ Players at (3, 4), (2, 4), and (1, 4) are meanwhile lulled into security by the large contingent in each treatment who do not challenge the previous high bid. By (2, 4), frontrunners let down their guard to the point that 40% or more actually lower their bid. The stage is thus set for a sly laggard to jump in and win a battle at a modest bid.

Across the *interior states*—those in which both players have at least one win to their name—competition is best summarized by escalation. In all three treatments, and regardless of whether these states are reached by winning or losing, players match or beat the previous high bid 67% to 92% of the time.²⁷ Reaching an interior state

²⁶The difference at (4, 1) between Win 300 Lose 0 and each of the other two treatments is statistically verifiable below the 1% significance level with a standard proportions test.

²⁷Two noted exceptions in Table 2 are arriving at (3, 1) by losing and at (1, 3) by winning. These lower propensities may reflect discouragement at the former and overconfidence at the latter.

Table 3: Fixed Effects Regression for Predicting a Player’s Bid at an Interior State

	L285	L150	L0
<i>Lost last round:</i>			
Winning Bid Last Round	1.13 ^{***} (0.060)	0.98 ^{***} (0.117)	0.78 ^{***} (0.112)
Abs. Bid Diff. Last Round	-0.79 ^{***} (0.129)	-0.27 [*] (0.117)	-0.53 [*] (0.246)
R^2	0.518	0.364	0.348
<i>Won last round:</i>			
Winning Bid Last Round	1.16 ^{***} (0.081)	1.13 ^{***} (0.101)	0.93 ^{***} (0.079)
Abs. Bid Diff. Last Round	-0.51 ^{***} (0.122)	-0.56 ^{**} (0.177)	-0.80 ^{**} (0.207)
R^2	0.670	0.589	0.330

Significance levels: * 10%, ** 5%, *** 1%. Cluster-robust std. errors at session level in parentheses.

is a mutual signal of each player’s engagement in the competition. The high degree of correlation between the previous winning bid and a player’s subsequent bid is borne out further by the fixed effects models in Table 3. Grouping bids at interior states by whether the player won or lost the last round, these models predict a player’s bid solely from the previous winning bid and the absolute difference in the two players’ bids last round.²⁸ Across treatments, increases in the previous winning bid are matched nearly one-for-one in the subsequent bid—although the effect is slightly larger following a win and with losing penalties.

Our final two hypotheses address the size of the winning margin (H4) and the importance of winning the initial battle of the tournament (H5). A natural consequence of last stand behavior is that the winning margin tends to decrease in the losing penalty. Landslide victories, on the other hand, frequently occur when players are prone to surrender. Table 4 reports the distribution of winning margins for completed tournaments within each treatment.²⁹ Differences across treatments are most noticeable with winning margins of either one or four battles. In support of our hypothesis, landslide victories are clearly more pronounced in the Win 300 Lose 0 treatment and account for over 40% of completed tournaments, 11 percentage points above the other two treatments. Conversely, the Win 300 Lose 0 treatment has the lowest occurrence of neck-and-neck victories.³⁰

²⁸Time in the fixed effects model is measured by both the tournament number and the round within the tournament. Fixed effects are included for each individual subject.

²⁹Counts underlying the percentages in Table 4 are adjusted to account for attrition from early termination. That is, observations for a winning margin of k are multiplied by $1/(0.9^{4-k})$.

³⁰The claim that the Win 300 Lose 0 treatment has both a higher frequency of landslide victories

Table 4: Winning Margins by Treatment (in %)

		Winning Margin			
		4	3	2	1
Win 15	Lose 285	28.2	25.0	25.2	21.6
Win 150	Lose 150	29.4	27.4	22.4	20.9
Win 300	Lose 0	40.7	25.4	21.1	12.8

Table 5: Percent who Ultimately Win the Tournament After Winning at a Tied Position

	After Win at (4,4)	After Win at (3,3)	After Win at (2,2)
L285	73.7 (0.008, 232)	81.6 (0.010, 87)	77.8 (0.665, 63)
L150	78.8 (0.000, 241)	80.7 (0.062, 57)	73.8 (0.859, 42)
L0	77.6 (0.000, 243)	71.8 (0.611, 71)	80.4 (0.496, 46)
50/50 win rate	65.625	68.75	75.0

Parentheses contain p-values and the number of observations from a two-tailed binomial test with a 50/50 win rate at each battle as the null hypothesis.

Initial battles are often viewed as pivotal in deciding the final outcome of a dynamic contest.³¹ Our theory predicts that even with a last stand, the defense of the lead will cause upsets to fall to a mere trickle. At a minimum, if the initial battle carries any special weight, it should predict tournament outcomes better than a coin toss at each battle. Table 5 shows the correlation between winning a battle when the tournament is tied and winning the entire tournament. In all treatments, a win at (4, 4) carries more predictive weight (by 8 to 13 percentage points) than determining the outcome of each successive battle by a coin toss. That predictive edge, however, is gone once the tournament reaches the tied position of (3, 3) in the Win 300 Lose 0 treatment. And it is gone from all treatments by (2, 2). The loss of a predictive edge further nullifies our defense of the lead hypothesis, but it adds credence to the pattern of escalation.³²

and a lower frequency of neck-and-neck victories than the other two treatments can be verified at the 1% significance level with a standard proportions test.

³¹For example, Klumpp and Polborn (2006) apply a dynamic contest model to the US presidential primaries with an eye toward the large campaign and media attention that Iowa and New Hampshire receive as the first states to vote.

³²Table 5 excludes tournaments that ended prematurely. However, the results do not substantially change if we instead tally which player is in the lead at the end of a tournament (whether it ended prematurely or after a player won four battles). Although the p-value for the Win 150 Lose 150 treatment after winning at (3, 3) does increase to 0.150 (76.5% with 81 observations).

5. Individual Behavior and Player Types

The dynamic patterns of competition we have identified are further illuminated by examining the full path of individual tournaments. Although these paths do not contain the full set of contingent responses necessary to construct a player’s behavior strategy, we are able to study each player’s *realized strategy*—that is, the set of history-contingent actions taken by the player along the realized path of play, as generated by the players’ behavior strategies. Players tend to adopt behavior strategies that lead to one of several distinctive categories of realized strategies, and a given player’s observed behavior often conforms to the same category repeatedly.

We identify four major realized strategies that arise frequently and repeatedly across the different experimental sessions: escalate when challenged, last stand, maximin, and passive when challenged. Escalation captures the notion of a bidding war, by far the most common behavior. The last stand realized strategy is characterized by sudden and fierce competition when a player falls behind by two or three states. The final two realized strategies distinguish two distinct paths to a surrender—whether it is from the start of the tournament (maximin) or when the competition mounts (passive when challenged). The former is a maximin strategy in the sense that, while anticipating a tournament loss, players maximize their payoff by bidding zero. A prominent nuance to these strategies, especially with regards to escalating, is the level of a player’s bid at the initial state (4, 4). Initial bids serve as a signal to some degree and most players vary it only slightly from tournament to tournament.³³ A high initial bid typically signals a fairly aggressive player; low initial bids, on the other hand, carry less information as they are commonly used by both aggressive and non-aggressive players. To capture this, we divide players not only by the category of realized strategy they predominantly use, but also by whether their initial bids tend to be low, moderate, or high.³⁴ We present our formal taxonomy below:

Taxonomy. *A player’s realized strategy in a single best-of-seven tournament is classified as follows:*

1. Escalate when Challenged: *In at least two battles on the realized path of the tournament the player submits a bid that is strictly greater than his own last*

³³For the last ten tournaments, the 25th, 50th, and 75th percentiles of the standard deviation of initial bids by player are 0.04, 2.10, and 4.89. The consistency of initial bids is likely due, at least in part, to the fact that players are matched at random for each tournament.

³⁴We define high bids as greater than 12 and low bids as less than 2. These cutoffs correspond to the 75th and 39th percentiles of initial bids across all sessions. Some initial bids are quite common: 0, 1, 2, 5, and 10 account for 54% of all initial bids, with 0 alone counting for 23%.

*bid and weakly greater than his opponent's last bid. The first occurrence may not be at a state where the player has already accrued two or more losses and the second occurrence may not be at (4, 1).*³⁵

2. Maximin: *The realized path of the tournament reaches either $(i, 1)$ or $(1, j)$, $i, j \in \{1, 2, 3, 4\}$, and the player does not place a bid in any battle that is both strictly greater than his own last bid and weakly greater than his opponent's last bid. Moreover, all of the player's bids must be strictly less than two.*
3. Last Stand: *The conditions for "escalate when challenged" are not satisfied, but the following properties are: Upon reaching $(4, 2)$, $(4, 1)$, or $(3, 1)$, the player bids strictly greater than his own last bid and either (i) weakly greater than his opponent's last bid; or (ii) strictly greater than the midpoint between his own last bid and his opponent's last bid but not less than two.*
4. Passive when Challenged: *The realized path of the tournament reaches either $(i, 1)$ or $(1, j)$, for $i, j \in \{1, 2, 3, 4\}$; the conditions for "escalate when challenged", "last stand", or "maximin" are not satisfied; and in any state in which the player has already incurred at least one loss the player bids strictly below his opponent's previous bid.*

A player is classified by one of these four categories of realized strategies if, in the last ten tournaments of play, the player's behavior conforms to that category at least three times and in a strict majority of the tournaments for which one of the above classifications applies. A player is further classified as a low, middle, or high initial bidder if a strict majority of bids at $(4, 4)$ are within the range of $0 - 1.9$, $2 - 12$, or greater than 12, respectively.

Roughly 66% of the 2160 person-level tournament observations over the last ten tournaments can be classified by one of these four categories of realized strategies. The remaining 34% is largely composed of tournaments that ended before a realized strategy could be distinguished.³⁶ Of the 216 players, 201 can be classified by their initial bid, 180 by a realized strategy, and 169 by both.³⁷

³⁵These exclusions are meant to distinguish escalate when challenged from a last stand.

³⁶A tournament needed to continue for at least three battles to make any classification. In many cases, four or more battles were needed. A total of 18% of the tournaments ended after the first or second battle and are therefore unclassified; an additional 4% ended after the third battle without being classified. Another large group of unclassified tournaments are those in which players faced non-aggressive rivals and consequently did not clearly demonstrate one of the categories of realized strategies. Specifically, 7% of the tournaments reached $(0, 4)$ without being classified.

³⁷There are two prominent reasons for a player not being classified by a realized strategy. The first is a lack of sufficient data due to either tournaments ending early, or players being paired

Table 6: Distribution of Player Types and Realized Strategies (in %)

	Player Type					Realized Strategy				
	MM	Pass	LS	Esc	Other	MM	Pass	LS	Esc	Other
L285	4.2	1.4	15.3	66.7	12.5	5.4	2.8	16.9	46.0	28.9
L150	4.2	1.4	15.3	62.5	16.7	6.3	2.2	17.8	40.3	33.5
L0	13.9	4.2	6.9	54.2	20.8	10.3	4.4	11.3	34.0	40.0

MM: Maximin; Pass: Passive when Challenged; LS: Last Stand; Esc: Escalate when Challenged; Other: all non-classified player types and realized strategies.

Table 6 summarizes both the distribution of player types and the distribution of realized strategies across the individual tournaments by treatment. As we have already noted, escalate when challenged is the most prevalent behavior. Between one-half and two-thirds of all players in all treatments fall under this type. Although the remaining player types are in the minority, their incidence across treatments is particularly telling. There is a relative abundance of maximin and passive when challenged players in the Win 300 Lose 0 treatment, a fact which reinforces our hypothesis of players surrendering when losing is costless. The opposite holds true for last stand players, which arise in greater frequency in the Win 15 Lose 285 and the Win 150 Lose 150 treatments. These differences can be substantiated statistically in terms of both player types and tournament-level realized strategies.³⁸

Just as players are classified in Table 6 by a majoritarian use of a realized strategy, we can also look at the actual number of times that they play a particular realized strategy. Players who stick exclusively to one realized strategy throughout the last

with non-aggressive rivals, in which case the characteristics of the different realized strategies are unlikely to be observed. Of the 36 players who were not classified by a realized strategy, 14 had five or more tournaments which could not be classified because the tournament either ended after one of the first three battles, or because the player was paired with a non-aggressive rival. The second common reason for remaining unclassified is having several tournaments under two or more categories of realized strategies, resulting in no clear majority.

³⁸Based on a one-tailed proportion test, maximin occurs more frequently in Win 300 Lose 0 than in Win 150 Lose 150 (p-value for player types: 0.021; p-value for realized strategies: 0.003). Likewise, maximin occurs more frequently in Win 300 Lose 0 than in Win 15 Lose 285 (p-value for player types: 0.021; p-value for realized strategies: < 0.001). Last stands are more common in Win 150 Lose 150 or Win 15 Lose 285 than in Win 300 Lose 0 (in either case, p-value for player types: 0.056; p-value for realized strategies: < 0.001). As a robustness check, Table A.1 in the appendix shows the distribution of realized strategies across treatments for several variations of the taxonomy definitions. The statistical evidence for last stands in the treatments with losing penalties and for maximin in the Win 300 Lose 0 treatment continues to hold across these variations.




Figure 3: CDF of number of times players use a strategy during last ten tournaments

ten tournaments of the experiment are in fact a minority with nearly 70% using two or more of the classifiable realized strategies at least once.³⁹ This practice of trying out the different strategies can be seen in Figure 3, which contains the cumulative distribution of the number of times players use a given realized strategy during the last ten tournaments. In every treatment, more than one-half of all players use the last stand realized strategy in one or more of these last ten tournaments. Usage rates for maximin and passive when challenged sit at roughly one-quarter, and nearly everyone has tried escalate when challenged. Notwithstanding this shopping around, usage rates for the different realized strategies are fairly delineated across treatments, with the distributions largely ranked in terms of stochastic dominance.⁴⁰

Use of the escalate when challenged strategy is crisply ordered across the three treatments. The Win 15 Lose 285 treatment has the highest occurrence at nearly

³⁹Over the last ten tournaments, 65 of the 216 players used only one realized strategy in the tournaments that can be classified; 101 used two, 45 used three, and 5 used all four.

⁴⁰Two such rankings are significant at the 10% level with a KS test: the last stand distribution is greater in the Win 150 Lose 150 than in the Win 300 Lose 0 treatment (p-value: 0.066); and escalate when challenged is greater in the Win 15 Lose 285 than in the Win 300 Lose 0 treatment (p-value: 0.044).

every level, followed by the Win 150 Lose 150 treatment, and then the Win 300 Lose 0 treatment.⁴¹ For maximin and passive when challenged, the Win 300 Lose 0 treatment comes close to first-order stochastically dominating the other treatments. In those other treatments, only a few players use maximin more than twice, and hardly any use passive when challenged more than once.⁴² The overall lack of observations for passive when challenged is partially due to two factors which are out of a player’s control: first, that the rival is aggressive; and second, that the tournament continues long enough for the player to demonstrate that he is adequately passive.

There is an interesting wrinkle in the ranking of the distributions for the last stand realized strategy. While the Win 300 Lose 0 treatment is unmistakably in third place, the ranking of the other two treatments switches midway. The likelihood of ever using the last stand realized strategy is highest in the Win 150 Lose 150 treatment, however, much of this likelihood is concentrated on using it only once or twice. By three times, there is no difference in the usage rate between the Win 150 Lose 150 and the Win 15 Lose 285 treatments. But the Win 15 Lose 285 treatment takes a slight lead thereafter. It is fitting that the last stand is employed so repeatedly in the treatment with the larger losing penalty.⁴³

We now turn to the question of how well each of these strategies performs. Figure 4 shows the average tournament payoff over the last ten tournaments of play for every player that can be classified by both a realized strategy and an initial bid. Here, players are coded by realized strategy, and escalate when challenged players are further coded by the level of their initial bid. Observations are also grouped in rows by experimental sessions. Maximin players serve as a natural baseline for comparison since all players could have guaranteed themselves this payoff. Factoring in the probability of a tournament ending early, consistently bidding zero leads to an average payoff of -207.8 in the Win 15 Lose 285 treatment, -109.4 in the Win 150 Lose 150 treatment, and 0 in the Win 300 Lose 0 treatment. As can be expected, maximin payoffs tend to cluster around these points.⁴⁴

⁴¹With a Mann-Whitney U test, the occurrence of escalation in the Win 15 Lose 285 is greater than in the Win 150 Lose 150 treatment (p-value: 0.060); it is likewise higher in the Win 150 Lose 150 than in the Win 300 Lose 0 treatment (p-value: 0.079).

⁴²Again by a Mann-Whitney U test, players use maximin at higher rates in Win 300 Lose 0 than in Win 15 Lose 285 (p-value: 0.092); passive when challenged in the Win 300 Lose 0 treatment also ranks above the Win 150 Lose 150 treatment (p-value: 0.086).

⁴³Although there is no overall statistical ranking between the Win 15 Lose 285 and the Win 150 Lose 150 treatments, last stand usage rates are significantly lower in the Win 300 Lose 0 treatment (Mann-Whitney U test p-values: 0.046 and 0.007, respectively).

⁴⁴The occasional maximin player would win a tournament—frequently by deviating from the maximin realized strategy when winning looked feasible (although there are a few scattered in-




Figure 4: Average payoff by player type and treatment. Observations from the same experimental session are grouped by row. Sessions are ordered by median payoff.

More than any other realized strategy, escalate when challenged with a high initial bid frequently fares worse than maximin.⁴⁵ Although escalation can be harmful if used too aggressively, it can also be profitable if used with an appropriate amount of moderation. Players who escalate with initial bids in the middle range often surface to the top of the payoff distributions. A mid-range initial bid provides an affordable starting point and frequently results in an early lead. Low-range initial bids, on the other hand, often forfeit the early lead.⁴⁶ Overall, last stand players do marginally better than maximin players.⁴⁷ Instead of simply swallowing the losing penalty, the last stand strategy seeks to buy the player time so that the tournament may end

stances where a tournament was actually won with a cumulative bid of zero).

⁴⁵The average payoff for the 41 players across treatments who escalated with a high initial bid was 10.9 below the expected payoff from maximin. However, with a Wilcoxon signed-rank test, these payoffs were not significantly below the maximin expected payoffs (p-value: 0.424).

⁴⁶Across treatments, average payoffs for the 54 mid-range and the 29 low-range escalation players are 53.2 and 37.5 above the maximin expected payoff. Each does significantly better than maximin (Wilcoxon signed-rank test p-values of 0.000). With a Mann-Whitney U test, however, the mid-range only weakly outperforms the low-range (p-value: 0.141).

⁴⁷Again grouping across treatments, the average payoff for the 24 last stand players is 12.5 above the expected payoff for maximin (p-value of 0.092 with a Wilcoxon signed-rank test).

early. Due to the lack of observations, passive when challenged is difficult to assess, but it appears to do quite well in the Win 300 Lose 0 treatment where two such players rank first in their sessions. These players are vested enough to win against easy competitors, but they are also quick to avoid costly confrontation.

6. Conclusion

Set against a backdrop of dynamic contests in the economic, political, and athletic landscapes, we use a controlled laboratory experiment to examine strategic patterns of competition in a best-of-seven tournament. Each battle within the tournament is modeled as an all-pay auction, and we vary the prize and penalty for winning or losing the tournament. Several specific insights come to light from this experiment. First, there is a persistent tendency across treatments to steadily escalate the cost of competition. After each player has won at least one battle, it is remarkably rare to see a winning bid decrease from one battle to the next. This is especially enlightening since escalation is either absent or plays at most a secondary role in many models of dynamic contests.⁴⁸ Behaviorally, bidding histories matter, and like the sunk cost fallacy, escalation can overemphasize the past relative to forward-looking Markovian strategies. Dismissing these histories would make Selten’s strategy method feasible, but even still, it may be that few dynamic contests are truly history independent.⁴⁹

Our principal hypotheses pertain to surrendering when losing is costless and making a last stand when losing has some bite. The clearest support for these come from studying behavior along the tournament path where one player consistently loses each battle. There, as utter defeat draws nigh, players are much more prone to spring into action when facing a sizable penalty—even if it entails a costly fight from that moment forward. The opposite is true when a penalty is absent. Looking at individual player behaviors, we find that there is a mixture of player types, including those who routinely surrender, make a last stand, or escalate. This mixture transforms the initial battles into a signaling game of sorts. Players who might otherwise escalate may be matched with someone who first appears to be passive, but later catches them off guard with a last stand. The key divisions then are whether the tournament is escalating, and if not, whether a player who appears to have surrendered has indeed surrendered.

⁴⁸Reaching back to the seminal work of Harris and Vickers (1987), a principal finding has been that early wins create momentum for the frontrunner and discourage the laggard. The emphasis therefore is on strategic positioning rather than escalation.

⁴⁹Malueg and Yates (2010) find that professional tennis matches do appear to exhibit state dependent Markov strategies (based on match outcomes rather than the effort players put forth).

Appendix A. Alternative Taxonomies

Table A.1: Distribution of Realized Strategies for Alternative Taxonomies (in %)

	L285	L150	L0	Proportion Tests ¹	
<i>Maximin:</i>					
Original	5.4	6.3	10.3	L0 > L285	L0 > L150
All bids less than one instead of two	4.4	6.1	10.1	0.000	0.003
<i>Passive:</i>					
Original	2.8	2.2	4.4		
<i>Last Stand:</i>					
Original	16.9	17.8	11.3	L285 > L0	L150 > L0
Meets (i) in original definition ²	16.1	17.5	10.7	0.001	0.000
Meets (ii) in original definition	15.7	17.4	11.1	0.001	0.000
Bid > sum of all own past bids ³	13.8	15.0	9.3	0.005	0.000
<i>Escalate:</i>					
Original	46.0	40.3	34.0		
Outbid ⁴ in at least three battles	34.0	26.7	23.9		
Outbid in all but possibly one battle ⁵	31.9	22.9	22.1		
Outbid in every battle	17.2	9.3	9.4		

Except for the noted changes, the alternative taxonomy definitions remain the same as the original definitions. Alternative definitions do not affect the distribution of the other realized strategies (when a realized strategy is partially defined as not being some other realized strategy, we assume that the original definition of that other realized strategy holds).

¹ P-values from one-tailed proportion tests. Alternative hypotheses are noted.

² In each last stand variation, the original conditions for escalate when challenged are not satisfied.

³ This bid, which is greater than the sum of the player's previous bids, must be at (4, 2), (4, 1), or (3, 1). It must also be weakly greater than the opponent's last bid.

⁴ Here, *outbid* means to place a bid strictly greater than your own last bid and weakly greater than your opponent's last bid.

⁵ The last two variations of escalate also require the tournament to continue for at least four battles.

Appendix B. Experiment Instructions

Thank you for your willingness to participate in this experiment. You will have the opportunity to earn some money as part of this experiment—the exact amount you earn will be based on both your choices and the choices of the other participants. Funding has been provided by the Economic Science Institute. You will be paid

privately at the conclusion of the experiment.

In order to preserve the experimental setting, we ask that you DO NOT talk with the other participants, make loud noises, or otherwise disturb those around you. You will be asked to leave and will not be paid if you violate this rule. Please raise your hand if you have any questions.

There are two parts to this experiment.

Part 1

In the first part of the experiment, you will be given a set of 15 choices. You will be asked to choose between receiving \$1 for sure (Option A) and receiving \$3 with some probability and nothing otherwise (Option B). The probability of winning \$3 in Option B varies across the 15 choices. You will receive payment for one of your choices. The computer will draw a number between 1 and 15 at random, and you will be paid for your choice corresponding to that number. If you chose Option A, you will receive \$1. If you selected Option B, the computer will randomly draw another number between 1 and 20, and the result of that draw will determine whether you are paid \$3 or \$0.

Are there any questions?

Part 2

The second part of the experiment consists of 20 best-of-7 tournaments. In each tournament, you will be paired at random on the computer with another participant. The winner of each tournament will receive a prize and the loser will incur a penalty.

The currency for this part of the experiment is rupees, and the exchange rate is 50 rupees = 1 US dollar. As part of this experiment you have received an account with 850 rupees (equivalent to \$17.00). This account is in addition to the \$7.00 show up fee. The prize for winning a tournament is 150 rupees, and the penalty for losing is 150 rupees.

In order to win a tournament you must be the first player to win 4 contests. A contest consists of entering a bid on the computer screen. The computer will allow bids that are either whole numbers or have up to one decimal point that are between 0 and 300 inclusive. You win a contest if your bid is higher than your opponent's (in case a tie occurs, the computer will decide the winner randomly, giving each player a 50% chance of winning). Once both players have entered their bids, the computer will display the two bids and indicate which player is the winner. The computer will also display past bids and the total number of contests that each player has won so far in the tournament.

After each contest, there is a 10% chance that the tournament will suddenly end.

The computer will randomly determine whether or not to end the tournament by selecting an integer between 1 and 10 (each number is equally likely to be drawn). If a 1, 2, 3, . . . , 9 is drawn, then the tournament will continue, and you will return to the bidding screen to bid in another contest. However, if the computer draws a 10, then the tournament will end early. Numbers that the computer has drawn previously may be drawn again. Given that no player has won 4 contests, there is always a 90% chance of continuing to the next round of the tournament. The following table shows the percent of all tournaments that are expected to reach a given round provided that no player has won 4 contests by that round.

Round	1	2	3	4	5	6	7
% of Tournaments	100%	90%	81%	73%	66%	59%	53%

Earnings

Your earnings for each tournament are based on your bids and whether you win or lose the tournament. All of your bids throughout the tournament will be subtracted from your earnings. Please note that each of your bids will be subtracted regardless of whether you win or lose each contest.

The prize of 150 rupees will be added to your earnings if you win the tournament, and the penalty of 150 will be subtracted from your earnings if you lose.

Here are some examples to illustrate how your earnings for a tournament are calculated. If you win a tournament in six rounds, then your earnings are as follows:

$$150 - (\text{Round 1 bid}) - (\text{Round 2 bid}) - (\text{Round 3 bid}) \\ - (\text{Round 4 bid}) - (\text{Round 5 bid}) - (\text{Round 6 bid})$$

Similarly, your earnings for losing a tournament in five rounds are given below:

$$-150 - (\text{Round 1 bid}) - (\text{Round 2 bid}) - (\text{Round 3 bid}) \\ - (\text{Round 4 bid}) - (\text{Round 5 bid})$$

If the computer does end the tournament before one of the players has won 4 contests, then neither player receives a prize or incurs a penalty. However, your bids will still be subtracted from your earnings. For example, if the computer stops the tournament after three rounds, you earn the following:

– (Round 1 bid) – (Round 2 bid) – (Round 3 bid)

When a tournament ends, either by a player winning 4 contests or by the computer ending it early, the computer will display your earnings for that tournament. You will then be paired at random with another participant for the next tournament.

We ask for your patience as there may be a short pause between tournaments. This may happen, for example, if your tournament ended early, but your next randomly selected partner is still competing in a tournament.

Payment

At the end of the experiment, 2 of the 20 best-of-seven tournaments will be selected at random. Your payment will be based on the average of your earnings in those 2 tournaments. The average will be added to your 850 rupee account and then converted from rupees to dollars (50 rupees = 1 US dollar). Positive earnings will increase the balance in your account, while negative earnings will decrease it. You will be paid the balance of your account.

Quiz #1

Your account initially has 850 rupees. The winning prize is 150, and the losing penalty is -150 .

Contest 1: Your bid: 45 Your opponent's bid: 73

Contest 2: Your bid: 92 Your opponent's bid: 100

Contest 3: Your bid: 21 Your opponent's bid: 21

Tournament randomly terminated after 3rd contest.

How many rupees would you receive from this portion of the experiment if this tournament was selected for payment?

Quiz #2

Your account initially has 850 rupees. The winning prize is 150, and the losing penalty is -150 .

Contest 1: Your bid: 295 Your opponent's bid: 23
Contest 2: Your bid: 70 Your opponent's bid: 150
Contest 3: Your bid: 51 Your opponent's bid: 40
Contest 4: Your bid: 80 Your opponent's bid: 20
Contest 5: Your bid: 72 Your opponent's bid: 80
Contest 6: Your bid: 51 Your opponent's bid: 70
Contest 7: Your bid: 200 Your opponent's bid: 175

How many rupees would you receive from this portion of the experiment if this tournament was selected for payment?

Quiz #3

Your account initially has 850 rupees. The winning prize is 150, and the losing penalty is -150 .

Contest 1: Your bid: 27 Your opponent's bid: 295
Contest 2: Your bid: 41 Your opponent's bid: 150
Contest 3: Your bid: 200 Your opponent's bid: 40
Contest 4: Your bid: 20 Your opponent's bid: 78
Contest 5: Your bid: 31 Your opponent's bid: 83

How many rupees would you receive from this portion of the experiment if this tournament was selected for payment?

This is the end of the instructions. If you have any questions, please raise your hand and a monitor will come by to answer them. If you are finished with the instructions, please click the Start button. The instructions will remain on your screen until the experiment begins. We need everyone to click the Start button before we can begin the experiment.

Subjects were given the following demographics survey at the end of the experiment.

1. What is your age? [Numeric Answer]
2. What is your major? [Free Response]
3. What is your gender? [(A) Male; (B) Female]
4. Have you previously participated in an economics experiment? [(A) Yes; (B) No]
5. On average, how many hours per week are you employed at a job? [(A) 31+ hrs/week; (B) 21–30 hrs/week; (C) 11–20 hrs/week; (D) 6–10 hrs/week; (E) 0–5 hrs/week]

6. How many hours per week do you spend studying outside of class? [(A) 31+ hrs/week; (B) 21–30 hrs/week; (C) 11–20 hrs/week; (D) 6–10 hrs/week; (E) 0–5 hrs/week]
7. Do you participate in club or intercollegiate athletics? [(A) Yes; (B) No]
8. How much time do you typically spend in student organizations or other extracurricular activities? [(A) 31+ hrs/week; (B) 21–30 hrs/week; (C) 11–20 hrs/week; (D) 6–10 hrs/week; (E) 0–5 hrs/week]
9. How regularly do you participate in a religious worship service? [(A) 1 or more times/week; (B) 1–3 times/month; (C) 1–3 times/semester; (D) 1–3 times/year; (E) Never]
10. Politically, do you consider yourself to be: [(A) Very liberal; (B) Somewhat liberal; (C) Neutral; (D) Somewhat conservative; (E) Very conservative]

Submission

Bid

1.0

The maximum bid is 300.

Submit Bid

Round Results

Round	My Bid	Other's Bid	Sum of My Bids	Result
1	23.9	15.1	23.9	win
2	28.1	5	52	win
3	22	26.2	74	loss
4	8.4	13.4	82.4	loss
5	1	28	83.4	loss
6				
7				

The person with the highest bid wins the round.
Ties are broken randomly.

The first person to win 4 rounds wins the tournament.

There is a 10% chance the tournament will end early after this round.

All of your bids will be subtracted from your earnings.
The winner of the tournament will receive 150; the loser must pay 150.
The average of your earnings from two randomly selected tournaments will be added to or subtracted from your account of 850.

Submit your bid.

Figure B.1: Bidding screen during a best-of-seven tournament

Acknowledgments

We express gratitude to Jeffrey Kirchner for programming the experiment and gratefully acknowledge helpful comments from Gabriele Camera, Brice Corgnet, Cary Deck, David Porter, Stephen Rassenti, Roman Sheremeta, Vernon Smith, Nathaniel Wilcox, Bart Wilson, John Wooders, and anonymous reviewers. This work has also benefited from participants at several conferences: 14th APET Conf. (Seattle, WA); 25th Int'l Conf. on Game Theory (Stony Brook, NY); 15th SAET Conf. (Cambridge, England); and seminars at the Economic Science Institute (Orange, CA) and the Max Planck Institute for Tax Law and Public Finance (Munich, Germany).

References

- Agastya, M., McAfee R.P., 2006. Continuing wars of attrition. Working Paper.
- Baye, M.R., Kovenock D., de Vries C.G., 1996. The all-pay auction with complete information. *Econ. Theory* 8(2), 291–305.
- Crawford, V.P., Costa-Gomes, M.A., Iriberry, N., 2013. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *J. Econ. Lit.* 51(1), 5–62.
- Dechenaux, E., Kovenock D., Sheremeta R.M., 2015. A survey of experimental research on contests, all-pay auctions and tournaments. *Exper. Econ.* 18(4), 609–669.
- Deck, C., Sheremeta, R.M., 2012. Fight or flight? Defending against sequential attacks in the game of siege. *J. Conflict Resolution* 56(6), 1069–1088.
- Duffy, J., 2008. Experimental Macroeconomics. In: Durlauf, S.N., Blume, L.E. (Eds.), *The New Palgrave Dictionary of Economics* (2e).
- Embrey, M., Fréchette, G.R., Lehrer, S.F., 2015. Bargaining and reputation: An experiment on bargaining in the presence of behavioral types. *Rev. Econ. Stud.* 82(2), 608–631.
- Ernst, C., Thöni, C., 2013. Bimodal bidding in experimental all-pay auctions. *Games* 4, 608–623.
- Fragiadakis, D.E., Knoepfle, D.T., Niederle, M., 2013. Identifying predictable players: Relating behavioral types and subjects with deterministic rules. Working Paper.
- Fudenberg, D., Gilbert, R., Stiglitz, J., Tirole, J. 1983. Preemption, leapfrogging and competition in patent races. *Euro. Econ. Rev.* 22(1), 3–31.
- Gelder, A., 2014. From Custer to Thermopylae: Last stand behavior in multi-stage contests. *Games Econ. Behav.* 87, 442–466.
- Gneezy, U., Smorodinsky, R. 2006. All-pay auctions—An experimental study. *J. Econ. Behav. Org.* 61(2), 255–275.
- Harris, C., Vickers, J., 1987. Racing with uncertainty. *Rev. Econ. Stud.* 54(1), 1–21.
- Hillman, A.L., Riley, J.G., 1989. Politically contestable rents and transfers. *Econ. Politics* 1(1), 17–39.
- Holt, C.A., Laury, S.K., 2002. Risk Aversion and Incentive Effects. *Amer. Econ. Rev.* 91(5), 1644–1655.
- Irfanoglu, Z.B., Mago, S.D., Sheremeta, R.M., 2014. The New Hampshire Effect: Behavior in Sequential and Simultaneous Election Contests. Working Paper.

- Klump, T., Polborn, M.K., 2006. Primaries and the New Hampshire effect. *J. Public Econ.* 90(6–7), 1073–1114.
- Konrad, K.A., Kovenock, D., 2005. Equilibrium and efficiency in the tug-of-war. CESIFO Working Paper No. 1564.
- Konrad, K.A., Kovenock, D., 2009. Multi-battle contests. *Games Econ. Behav.* 66(1), 256–274.
- Mago, S.D., Sheremeta, R.M., 2012. Multi-battle contests: An experimental study. Working Paper.
- Mago, S.D., Sheremeta, R.M., Yates, A., 2013. Best-of-three contest experiments: Strategic versus psychological momentum. *Int. J. Indust. Org.* 31(3), 287–296.
- Malueg, D.A. and Yates, A.J., 2010. Testing contest theory: Evidence from best-of-three tennis matches. *Rev. Econ. Stat.* 92(3), 689–692.
- Noussair, C., Matheny, K., 2000. An experimental study of decisions in dynamic optimization problems. *Econ. Theory* 15(2), 389–419.
- Potters, J., De Vries, C.G., Van Winden, F. 1998. An experimental examination of rational rent-seeking. *Euro. J. Political Econ.* 14(4), 783–800.
- Selten, R. (1967). Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperiments. In: Sauermann, H. (Ed.), *Beiträge zur experimentellen wirtschaftsforschung*, 136–168. Tübingen: Mohr.
- Sheremeta, R.M., 2010. Experimental comparison of multi-stage and one-stage contests. *Games Econ. Behav.* 68(2), 731–747.
- Sheremeta, R.M., 2013. Overbidding and heterogeneous behavior in contest experiments. *J. Econ. Surveys*, 27, 491–514.
- Zizzo, D.J., 2002. Racing with uncertainty: A patent race experiment. *Int. J. Indust. Org.* 20(6), 877–902.