

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. Sports competitions involve the cumulative performance of a series of interactions over the course of a race, game, or match. Businesses may spend weeks or months jockeying for a lucrative contract. Political campaigns frequently span a year or more and major wars span several.

[☆]An 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)

While the time horizon may vary, a unifying feature of many contests is that the overall contest is often comprised of several smaller battles or component contests. Moreover, as participants compete in these battles they have some awareness of how they fare relative to their rivals—whether they are approaching victory or falling behind. Competition typically requires participants to make costly outlays, such as effort, money, or troops, so a chief strategic consideration is the amount and timing of these expenditures. Because a participant’s proximity to victory or defeat may determine the efficacy of these expenditures, allocations and outcomes in previous battles may strongly influence behavior later in the contest.

A variety of strategies may ensue in contests that are fought over a sequence of battles. Participants may fight intensely in each battle, especially when challenged, essentially escalating the conflict. Or they may only fight aggressively at selective key battles. A last stand, for instance, is classically characterized by fierce competition along the brink of an overall loss. Some participants may instead opt to forgo the cost of competing by folding when the competition gets tough or even from the get-go, virtually surrendering. In comparing and selecting among strategies in dynamic multi-battle contests, participants may attempt to achieve favorable outcomes by assessing the potential costs and rewards of the various combinations of actions that they and their rivals can feasibly take, and by trying to anticipate rival behavior. Standard tools in the formal analysis of optimal decision making in such contests are the theory of dynamic games and laboratory experiments. The purpose of this article is to examine a specific form of a multi-battle contest experimentally in a laboratory setting, test the major game theoretic predictions of behavior in that contest, and glean whatever insights into dynamic behavior that we can extract from the data.

The contest we examine is a two-player best-of-seven tournament under complete information. The tournament is comprised of sequentially-played component contests that are modeled as all-pay auctions. In an all-pay auction the high bidder wins, but both players incur 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 is the tournament winner.¹ The tournament winner receives a positive winning prize and the loser pays a nonnegative losing penalty.

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.

Konrad and Kovenock (2009) theoretically examine the set of subgame perfect equilibria in a class of contests that includes this type of contest under the assumptions that the winner of the tournament receives a positive prize, the loser pays a zero penalty, and there is no time discounting of the prize winnings or bid costs accrued in the course of competing in the individual all-pay auctions.² Konrad and Kovenock find that, when symmetrically situated, players compete intensely, but that a player who is behind will completely give up, unless there is a separate intermediate prize for winning each of the individual battles.³

This work was extended by Gelder (2014) to include strictly positive loser penalties, as well as the discounting of potential future earnings or losses over the component contests. In the absence of a loser penalty, time discounting is innocuous; in the unique subgame perfect equilibrium a player who is behind still completely gives up, unless there are intermediate prizes for winning individual battles. However, Gelder (2014) shows that, under certain conditions, the inclusion of loser penalties and time discounting induces a player who is close to losing the tournament to exhibit last stand behavior: the player bids more aggressively than his rival even though his rival is much closer to winning the tournament. Last stand behavior arises because, with discounting, if the penalty for losing the tournament is sufficiently large, the value of delaying that penalty for at least one more period may outweigh the rival's gain from securing an immediate win. In contrast, if a player is behind in the tournament by only one component contest, it is no longer simply a question of delaying defeat but also of potentially overtaking the lead. Since this poses a serious threat to the rival, the rival is predicted to bid much more aggressively in these states than the player who is behind.

This article experimentally investigates the theoretical predictions of Konrad and Kovenock (2009) and Gelder (2014). We examine best-of-seven tournaments under three prize-penalty combinations: one in which there is no penalty, one with a prize and penalty of equal magnitude, and a third in which the penalty is dramatically larger than the prize. For each of these three cases, the net difference between the prize and penalty is identical; so in a one-shot contest with risk neutral players,

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.

³An early example of a dynamic contest where players slacken their effort or even give up entirely if they fall behind is Fudenberg et al. (1983), who examine preemption in patent races by firms who have a marginal lead over their competitors.

equilibrium behavior would be the same in each case. These values were selected to allow for a test of the prediction of last stand behavior in a multi-battle contest with a winning prize and sufficiently large losing penalty, as well as the contrasting prediction of the surrender of a player who falls behind in the corresponding contest with no penalty.

The diverging predictions of surrendering versus making a last stand only come into play in the dynamic setting. Because discounting is a central ingredient for the last stand predictions of Gelder (2014), we implemented discounting within the experiment via a probability that the tournament would suddenly end after each component contest. If a tournament ended before either player had won four component contests, then neither player would receive the winning prize or pay the losing penalty, but both players would still be responsible for the payment of their previous bids.⁴

Although previous experimental work has addressed best-of-three tournaments (see Sheremeta 2010; Mago and Sheremeta 2012; Mago et al. 2013; and Irfanoglu et al. 2014), the source of the contrasting dynamics of last stand and surrendering behavior can only be realized with a larger number of battles. The best-of-seven tournament provides a horizon that is long enough to capture the desired dynamics, but short enough to keep the experiment simple. It is also a natural choice, as it is used in many sports championship settings, such as Major League Baseball’s World Series, the National Basketball Association’s Finals, and the National Hockey League’s Stanley Cup Final.

Players appear to gravitate toward different strategies and repeatedly use those strategies throughout the experiment. In this regard, we find that 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,⁵ we refer to these observable actions as a player’s *realized strategy*. We find that the same core set of realized strategies independently appear time and time again across the different experimental sessions. The most common is to engage in a bidding war, a realized strategy which we refer to as *escalate when challenged*. Our hypothesized surrendering behavior is reflected by two distinct realized strategies, *maximin*

⁴Under the assumption of risk neutrality, this adaptation of the tournament is theoretically innocuous.

⁵Because of the complexity of the tournament’s extensive form, the strategy method (Selten 1967) appears difficult to implement without strong and unrealistic restrictions on the information available to players at each stage of the game.

and *passive when challenged*. We also observe a remarkably pronounced *last stand* realized strategy. Defining a formal taxonomy, we classify the occurrence of each realized strategy and find that the surrendering strategies are particularly prominent in the treatment with no losing penalty. These same surrendering strategies are conspicuously scarce in the other two treatments that do have losing penalties. Conversely, the last stand realized strategy is a salient feature of the treatments with losing penalties but is less common when a losing penalty is absent.

Winning margins in the overall tournament are naturally affected by the choice of player strategies. Theory predicts that neck-and-neck outcomes are more likely to be observed when the losing penalty is relatively large, while landslide victories are more probable when the winning prize is dominant. We find evidence to support this hypothesis. On the other hand, theory predicts that under all three prize-penalty combinations a player who has a one battle lead over his rival in the tournament will compete more aggressively than his opponent. The aggregate data do not support this hypothesis.

Although the winning prize and losing penalty in each treatment are fixed, allowing players to expend or conserve resources through the size of their bids makes these tournaments non-constant-sum games—a fact which is further compounded by the possibility that these tournaments may suddenly terminate after any component contest with no winner and no loser. Given the strategic complexity of this environment, we examine two separate subject pools: one in which all subjects had previously participated in some separate contest related experiment, and another comprised of a mix of subjects with and without such prior experience. Although there are a few noted exceptions, the two subject pools behave quite similarly.

Our paper fits within a small but emerging literature on dynamic contest experiments, as well as within a broader literature on contests and tournaments (see Dechenaux et al. 2014 for an extensive survey on experiments involving contests). In terms of “best-of” experiments, we bridge the gap between the work on best-of-three tournaments mentioned previously and the best-of-19 tournament in Zizzo (2002), which was explicitly patterned after Harris and Vickers (1987) and used a logit-style lottery contest success function. Our paper is also closely related to the game of siege experiment by Deck and Sheremeta (2012). In their experiment, players are positioned asymmetrically so that one player (the defender) needs to win two successive battles to be victorious, while the attacker only needs to win one (this is the dynamic counterpart of a weakest link contest). The asymmetric starting point assumed by Deck and Sheremeta can be reached as an intermediate stage in a best-of-three and a best-of-seven tournament.

This paper also fits within the behavioral literature on heterogeneous player types in experiments. It can be expected that subjects will approach experimental settings from varying degrees of strategic sophistication, with the spectrum ranging from mere guesses to deliberate best-responses. *Level-k theory*, for instance, has been developed to classify and better understand this range of behavior.⁶ Other experimental work seeks to induce player types by allowing subjects to potentially be matched against a computer which is known to be programmed a certain way (see Embrey et al. 2014). Here, instead of inducing player types or focusing on the degree to which players are best-responding, we analyze the frequency and interaction of a handful of player types that arise endogenously. We also provide a rough ranking of these types in terms of average payoffs. To our knowledge, this is the first dynamic contest experiment to examine multiple player types.

We begin by giving a brief description of the theoretical framework in Section 2 and then describe how we set up the experiment in Section 3. Our analysis is in two parts. First, Section 4 provides a baseline analysis of our hypotheses within the aggregate data. We then address the phenomenon of multiple player types in Section 5.

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 can 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 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

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 largely be captured by such theories, more needs to be done in modeling play that is non-strategic but 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 unidimensional state space. Within the tug-of-war setting, 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.

⁸Although an arbitrary tie-breaking rule typically suffices, the equilibrium in Konrad and Kovenock (2009) requires that ties be awarded to the player who is ahead in the tournament.

	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

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:

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

Given these distributions, the expected payoffs are $u_H = \zeta_H - \zeta_L$ and $u_L = 0$; the 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), as well as in Gelder (2014), is in extending the one-shot all-pay auction to a dynamic structure where an actual prize is awarded only after a player has achieved a critical number of wins. Hence, it becomes necessary to identify the prize valuations at each interior state (i, j) where $i, j > 0$. 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

This assumption allows the frontrunner to coast to victory with a bid of zero when the laggard surrenders. Since this is a rather technical requirement, we use a fair randomizing device to break ties in the experiment.

there is no incentive to compete.⁹ In Gelder’s framework on the other hand, 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 that players actively compete at every interior state.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 degree of 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. The likelihood of the underdog catching up after an unsuccessful last stand is minimal at best. 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 it entails a much higher expenditure from the frontrunner than from the underdog in expectation. Thus the last stand acts as a defensive push, while the defense of the lead acts as an offensive one.

Based on the theoretical predictions, there are five main hypotheses that we will examine in this experiment. The first three address in turn the last stand, the tendency to surrender, and the defense of the lead. The fourth examines winning margins, a feature which is 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 overall outcome of the tournament.

⁹Since the player who is behind receives zero from continuing to lose, and since the expected payoff from winning a single state is also zero, then the prize valuation is zero as well. Konrad and Kovenock also examine the case where there is an intermediate prize for winning each battle. In that setting, the prize valuation for a player who is behind is solely based on the intermediate prize.

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.** *If there is no losing penalty, then a player who is behind will surrender (or cease to compete).*
- H3.** *Players with a one battle lead in the tournament will compete more aggressively than their opponent in order to maintain their lead.*
- H4.** *The expected winning margin is increasing in the size of the winning prize relative to the losing penalty.*
- H5.** *The winner of the initial battle will win the tournament the majority of the time.*

3. Methodology

We conducted 18 experimental sessions, each composed of 12 subjects. These sessions were conducted at the Economic Science Institute, Chapman University, in computer labs where the computers were separated by partitions for privacy. The experiment began with subjects reading the instructions on their computer (a copy of the instructions is provided in the appendix). After reading the instructions, subjects were given a short quiz comprised of three possible scenarios for how a best-of-seven tournament could unfold. Subjects were then asked to compute the payoff for each scenario. The purpose of this short quiz was to ensure that subjects had a basic level of comprehension about the structure of the game. The quiz was immediately followed by a short risk preference lottery *à la* Holt and Laury (2002). During the main portion of the experiment, subjects participated in 20 best-of-seven tournaments. Subjects were randomly and blindly paired and re-paired for each of these tournaments via the computer network. 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 of a best-of-seven tournament was treated as an all-pay auction: subjects placed bids simultaneously and the high bidder won (ties were broken randomly). In the all-pay fashion, the sum of a player's bids throughout a tournament was deducted from his or her payoff for that tournament. Additionally, the winner of the tournament received a prize and the loser incurred a penalty. Since time preferences for winning or losing in Gelder (2014) were implemented through a discount factor, and since discounting is difficult to replicate in a short laboratory experiment, we followed a common practice from macroeconomic experiments by implementing discounting via a continuation probability (see, for instance, Duffy 2008, and Noussair

and Matheny 2000). Until a player had succeeded in winning four battles, there was a 90% probability that the tournament would actually continue from one battle to the next (or in terms of 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. Our justification for this approach is that, under risk neutrality, a continuation probability is equivalent to discounting in terms of expected payoffs.¹¹

We conducted three separate payoff scenarios: the first 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).¹² Prizes, penalties, as well as all bids, were denominated in an experimental currency called rupees, where 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. The two treatments with non-zero penalties coincide with the Gelder (2014) model, while the treatment with no losing penalty fits the Konrad and Kovenock (2009) model. For each treatment, we ran a total of six experimental sessions—three of which were open to all individuals in our subject pool, and three were specifically limited to subjects who had previously participated in an experiment involving contests or contest theory. We will refer to the first subject pool as mixed and to the second as experienced.¹³ A summary of the experimental sessions by treatment is shown in Table 1. Bid observations in this table are limited to those during the last ten tournaments since that will be the focus of our analysis.

¹¹The random ending rule may also be thought of as the potential that some exogenous factor suddenly disrupts the conflict (such as the cavalry coming to save the day). An alternative method for implementing discounting is to make the size of the prize and the penalty contingent on the winning margin. Since the winning margin is based on the number of rounds in which players compete, this is a present value interpretation of discounting. A benefit of using the random ending rule in an experimental setting 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. They found similar results using each method.

¹²A player in the experiment who has won or lost four battles receives the prize or penalty with certainty. In the theoretical model, however, there is still discounting between the states $(i, 1)$ and $(i, 0)$ (as well as between $(1, j)$ and $(0, j)$). Hence, the experimental parameters must be multiplied by $1/\delta$ to coincide with the theoretical parameters.

¹³The Economic Science Institute at Chapman University conducts a large volume of economic experiments and keeps records of the different experiments subjects participate in. We made no attempt to regulate the number of subjects in the mixed sessions who had previously enrolled in a contest related experiment (the average was 6.4 with a minimum of three and a maximum of eleven).

Table 1: Sessions by Treatment

Sessions	Subject Pool	Prize	Penalty	Bid Observations
3	Experienced	15	285	1650
3	Experienced	150	150	1500
3	Experienced	300	0	1530
3	Mixed	15	285	1590
3	Mixed	150	150	1614
3	Mixed	300	0	1420

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 Figure 2. The bidding screen would also alert subjects when a tournament had finished, either by a player winning four rounds or by the computer ending the tournament early. After displaying the final outcome and payoffs for the tournament, subjects would then be randomly re-matched to begin a new tournament.

At the start of 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. At the end of the experiment, subjects received cash payment for the average balance of two randomly selected tournaments. Since losing penalties varied across treatments, and since bids and losing penalties were both 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 across treatments 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 to balance two opposing constraints: having the effective bidding budget be large enough so that players would not feel budget constrained, especially in tournaments that continued to the sixth or seventh battle; but also small enough so that the losing

¹⁴Since the unique subgame perfect equilibrium is also Markov perfect, equilibrium bids at the current state are not affected by the knowledge of bid realizations at previous states. However, in conducting the experiment we did not want to exogenously impose Markovian behavior on the subjects. One implication of this methodological choice was that the strategy method pioneered by Selten (1967) was not feasible due to the complexity of the game’s extensive form.

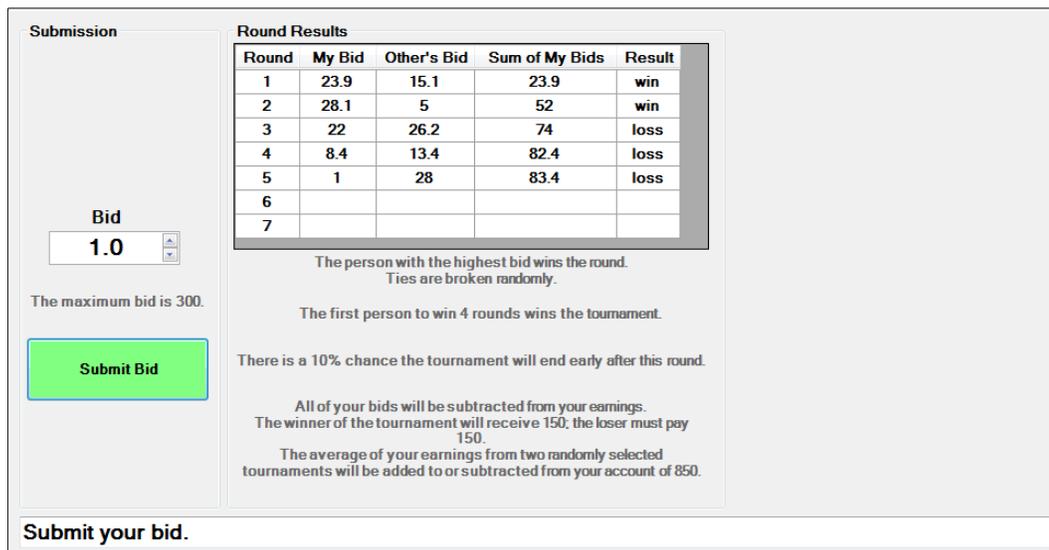


Figure 2: Bidding screen during a best-of-seven tournament

penalties would have some bite.¹⁵ For each round of a best-of-seven tournament, we allowed players to bid between 0 and 300 inclusive (with up to one decimal place).¹⁶

4. Initial Results

4.1. Summary Statistics

Before analyzing the main hypotheses, we briefly highlight the major summary statistics. We specifically summarize bidding observations, winning probabilities, and the distribution of bids by state and treatment. The fundamental level of observation is a player's bid at a particular state (i, j) within tournament t . As a whole, the data form a panel with twenty tournament observations per subject and up to seven bid observations per tournament. Variance in the bidding observations

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.

¹⁶The maximum bid of 300 corresponds to the upper bound of the equilibrium bidding distribution at state $(1, 1)$. The equilibrium distributions at all other states have upper bounds which are less than 300 (see Table 4).

Table 2: Number of Bidding Observations by State (i, j) and by Treatment

		Win 15 Lose 285				Win 150 Lose 150				Win 300 Lose 0			
Experienced	0	39	31	32	22	34	31	30	15	54	27	22	15
	1	54	56	58	44	55	63	45	30	66	47	40	30
	2	95	91	78		95	72	44		102	67	60	
	3	169	122			157	92			170	96		
	4	360				360				360			
Mixed	0	40	32	25	22	47	37	20	27	51	32	22	9
	1	66	54	51	44	77	61	51	54	59	41	33	18
	2	101	71	70		118	67	56		91	68	56	
	3	166	98			163	70			156	90		
	4	360				360				360			
Pooled	0	79	63	57	44	81	68	50	42	105	59	44	24
	1	120	110	109	88	132	124	96	84	125	88	73	48
	2	196	162	148		213	139	100		193	135	116	
	3	335	220			320	162			326	186		
	4	720				720				720			
		4	3	2	1	4	3	2	1	4	3	2	1

is considerably higher during the initial tournaments of the experiment since subjects are learning the structure of the game. Therefore, the analysis in this paper is solely based on the last ten tournaments.

Table 2 shows the number of bidding observations by treatment at state (i, j) where $i \geq j$. To reiterate, i is the number of battles that a player still needs to win in order to win the tournament, while j is the number of battles that the player's opponent still needs to win. The i index (4, 3, 2, 1) is shown at the bottom of the table, and the j index (4, 3, 2, 1, 0) is at the left. Due to the symmetry of the tournament, whenever one player is at (i, j) , their opponent is at (j, i) , so the table only shows states where a player is behind or the tournament is tied. The random ending rule causes the total number of observations to decrease by roughly 10% after each of

Table 3: Winning Percentages in State (i, j) : Theoretical and Observed

		Win 15 Lose 285				Win 150 Lose 150				Win 300 Lose 0			
Theory	1	56.4	52.4	5.2	50	26.3	26.3	2.6	50	0	0	0	50
	2	52.9	5.6	50		27.8	2.8	50		0	0	50	
	3	5.6	50			2.9	50			0	50		
	4	50				50				50			
Experienced	1	27.8	44.6	44.8	50	38.2	50.8	33.3	50	18.2	42.6	45.0	50
	2	36.8	45.1	50		35.8	33.3	50		22.5	46.3	50	
	3	40.2	50			34.4	50			31.8	50		
	4	50				50				50			
Mixed	1	39.4	40.7	51.0	50	39.0	39.3	60.8	50	13.6	22.0	33.3	50
	2	29.7	52.1	50		30.5	44.8	50		29.7	42.6	50	
	3	34.9	50			22.1	50			32.7	50		
	4	50				50				50			
Pooled	1	34.2	42.7	47.7	50	38.6	45.2	47.9	50	16.0	33.0	39.7	50
	2	33.2	48.1	50		32.9	38.8	50		25.9	44.4	50	
	3	37.6	50			28.1	50			32.2	50		
	4	50				50				50			
		4	3	2	1	4	3	2	1	4	3	2	1

the first four battles.¹⁷ In successive states, the number of observations continues to decrease through the random ending rule, but also decreases through players winning or losing tournaments.

Theoretical and observed probabilities of winning a battle at each state are shown in Table 3. Symmetry allows us to again focus on the states where a player is behind or the tournament is tied. The major patterns of competition can be seen

¹⁷For instance, in the pooled data of the Win 300 Lose 0 treatment, there are a total of 720 observations in the first round at $(4, 4)$. Of these, 90.6% persist to the second round—326 at $(4, 3)$ and an additional 326 at $(3, 4)$.

by examining the theoretical winning probabilities. For instance, the defense of the lead (H3) is reflected by the remote winning probabilities at states $(4, 3)$, $(3, 2)$, and $(2, 1)$ in the Win 15 Lose 285 and the Win 150 Lose 150 treatments. The last stand (H1) is evidenced by the underdog having the higher winning probability at states $(4, 2)$, $(4, 1)$, and $(3, 1)$ of the Win 15 Lose 285 treatment. Although not as strong, the underdog is still expected to win roughly a quarter of the time at these three states in the Win 150 Lose 150 treatment.¹⁸ Finally, the tendency to surrender (H2) is depicted in the Win 300 Lose 0 treatment by the zero probability of winning a battle when a player is behind in the tournament.

Although the observed probabilities from the laboratory fail to capture the defense of the lead, the basic contrast between making a last stand and surrendering can be seen. In the pooled data for instance, the winning probability at $(4, 1)$ falls to 16% in the Win 300 Lose 0 treatment—less than half the corresponding values of 34.2% and 38.6% in the two treatments with a losing penalty. In the mixed group, the Win 15 Lose 285 and Win 150 Lose 150 treatments actually have sizable jumps in the winning probability from $(4, 2)$ to $(4, 1)$ of 8 to 10 percentage points. While such a jump is absent in the experienced Win 15 Lose 285 treatment, it does boast the highest winning probabilities at $(4, 3)$ and $(4, 2)$, indicating that players were more likely to regain lost ground earlier in the tournament.

While the winning probabilities address the relative size of bids between (i, j) and (j, i) , it is also informative to have an absolute measure of bids at the different states, both theoretically and in the experiment. The subgame perfect equilibrium bidding distributions can be fully characterized with two sets of numbers: the size of the mass point at zero and the upper bound of the bidding distribution (see Equation 1; above the mass points, players uniformly randomize between zero and the upper bound). These are both presented in Table 4 by treatment and state. The mass points largely mirror the major features of the theoretical winning probabilities.¹⁹ A prominent feature of the upper bounds is that bids at states where the tournament is tied far and away exceed those at any other state—even when a losing penalty is present. We can also see that bids along the main diagonal of the tournament are increasing in the relative size of the winning prize, while competition off of the main diagonal is increasing in the size of the losing penalty.

¹⁸For the Win 150 Lose 150 treatment, a best-of-seven tournament is not large enough to include the states where the player who is behind wins battles with more than one-half probability.

¹⁹In the Win 300 Lose 0 treatment, the fact that both players always bid zero when one player is ahead is an artifact of the tie-breaking rule in Konrad and Kovenock (2009) that was mentioned previously.

Table 4: Theoretical Distributions: Mass Point at Zero (top row); Upper Bound (bottom row)

Win 15 Lose 285					Win 150 Lose 150					Win 300 Lose 0				
1	0	0	0.90	0	0.47	0.47	0.95	0	1.0	1.0	1.0	0		
2	0	0.89	0	0	0.44	0.94	0	0	1.0	1.0	0	1.0		
3	0.89	0	0	0.05	0.94	0	0	0	1.0	0	1.0	1.0		
4	0	0	0.06	0.13	0	0	0	0	0	1.0	1.0	1.0		
1	25.9	27.2	28.5	300.0	15.0	15.0	15.0	300.0	0	0	0	300.0		
2	22.0	24.4	244.4	28.5	13.5	13.5	256.5	15.0	0	0	270.0	0		
3	19.8	197.9	24.4	27.2	12.2	218.7	13.5	15.0	0	243.0	0	0		
4	160.3	19.8	22.0	25.9	185.9	12.2	13.5	15.0	218.7	0	0	0		
	4	3	2	1	4	3	2	1	4	3	2	1		

As a rule, bids within the experimental data are much lower than predicted when the tournament is tied, but can often be considerably higher than predicted at the other states. Figure 3 shows the empirical bidding distributions at each state in the pooled data; and for further detail, the 25th, 50th, and 75th percentile bids for each state and treatment are shown in Table 5. A pattern that is present in every experimental treatment, as well as in the theoretical distributions, is that bids progressively increase as the tournament proceeds to the top-right. It is uncommon for the median bid to be above ten in any treatment until both players have won at least one battle. Thereafter the median bids quickly rise as the tournament becomes more closely contended. By (1, 1), bids of 50 to 100 are commonplace. At states along the left edge of the tournament, at least a quarter of the bids are zero. These mass points are particularly conspicuous in the Win 300 Lose 0 treatment where the median bid at (4, 2) is zero in the pooled data, and even the 75th percentile bid is only 0.1 at (4, 1). Not only are these players surrendering at these states, but their rivals are responding with progressively lower bids. By (1, 4) in the experienced group, a mere bid of 1.0 marks the median. There is a different behavior in the Win 15 Lose 285 treatment—the primary difference being what happens at the top of the distribution. Now instead of 0.1, the 75th percentile bid at (4, 1) is 19 or 20. The mixed treatment is particularly suggestive of last stand behavior since the 75th percentile jumps from a bid of 10 at (4, 2) to a bid of 20 at (4, 1).

The shape of the bidding distributions is rather interesting in light of past experimental results. In one-shot contest settings, the experimental bidding distribution is frequently bifurcated with subjects submitting either high or low bids, but largely

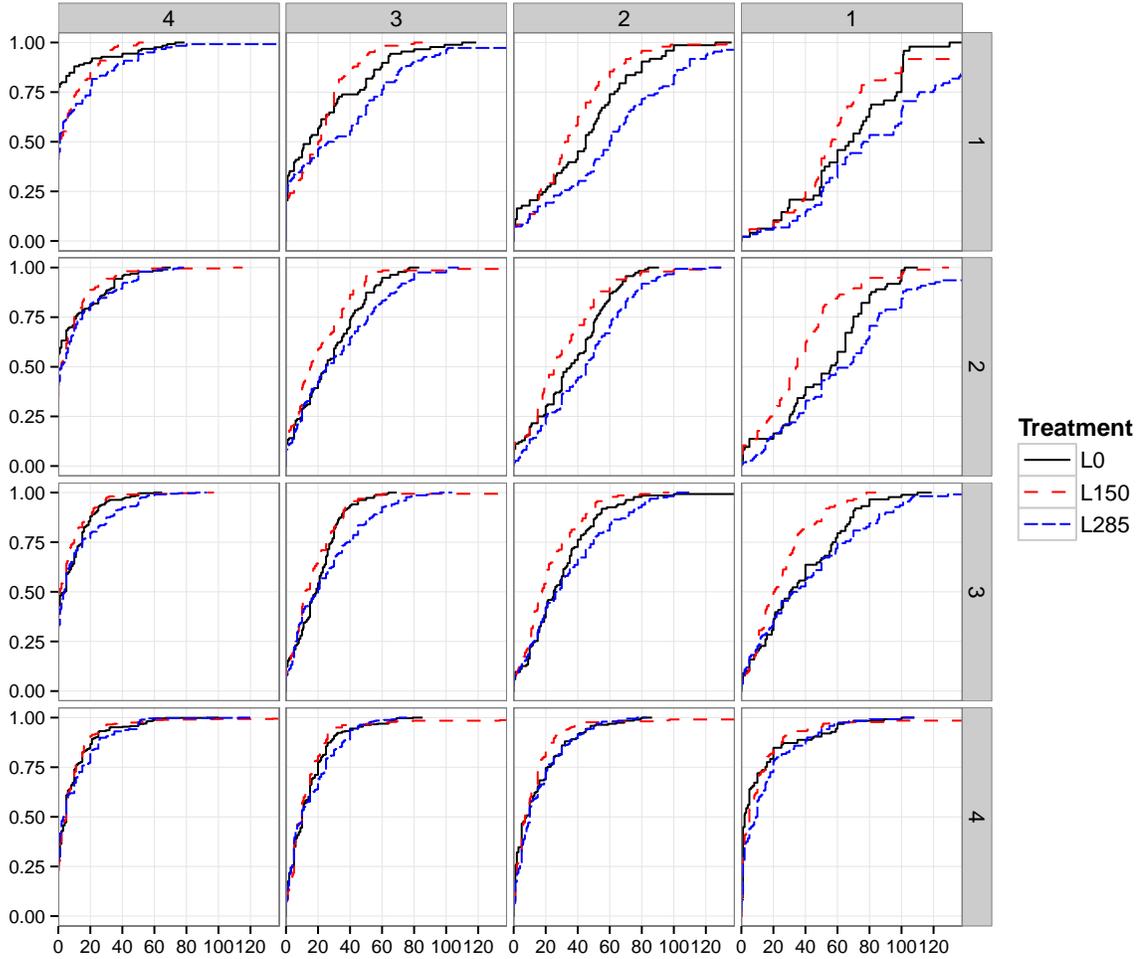


Figure 3: CDF of bids at each state by treatment in the pooled data

avoiding bids in the middle range.²⁰ Ernst and Thöni (2013) have demonstrated that prospect theory provides a possible explanation for this behavior.²¹ Here, however, instead of bifurcated distributions that are concave over the lower bidding range and convex over the upper range, bidding distributions at most states are concave throughout. Even at states (2, 1), (1, 1), and (1, 2) where the shape of the distribution appears to change a little, the curvature goes in the opposite direction with more mass being placed on intermediate bids—slightly convex in the lower bidding range and concave in the upper range.

²⁰See, for instance, Potters et al. (1998), Gneezy and Smorodinsky (2006), and Ernst and Thöni (2013).

²¹The prospect theory explanation has also been applied to Tullock contest experiments (see Sheremeta 2013).

Table 5: 25th, 50th, and 75th Percentiles of Bids at State (i, j)

		Win 15 Lose 285				Win 150 Lose 150				Win 300 Lose 0			
		<i>4</i>	<i>3</i>	<i>2</i>	<i>1</i>	<i>4</i>	<i>3</i>	<i>2</i>	<i>1</i>	<i>4</i>	<i>3</i>	<i>2</i>	<i>1</i>
Experienced	<i>1</i>	0.0	0.0	40.3	50.0	0.0	3.8	15.0	39.2	0.0	1.0	25.2	49.2
		1.0	28.0	67.5	98.0	3.0	20.0	26.0	50.0	0.0	15.0	46.0	66.0
		19.2	60.0	100.0	137.8	12.5	30.0	35.0	60.8	0.1	47.5	65.2	95.0
	<i>2</i>	0.0	15.0	30.0	38.2	0.0	7.0	7.8	15.0	0.0	5.0	14.0	29.0
		5.0	38.1	51.5	77.4	5.0	14.1	20.0	30.0	0.0	20.0	30.0	52.8
		20.0	60.0	70.0	89.0	10.0	25.0	30.4	40.0	5.0	40.0	54.8	70.8
	<i>3</i>	0.0	8.3	18.0	14.2	0.0	7.0	10.0	11.0	0.0	6.0	10.0	11.5
		6.0	25.5	35.0	40.0	5.0	11.0	16.0	16.0	1.0	16.0	20.0	26.0
		20.0	42.0	58.5	70.0	10.0	21.0	25.5	30.9	11.8	27.2	35.5	53.5
	<i>4</i>	0.1	5.0	5.0	1.0	0.1	5.0	2.0	1.8	0.0	2.0	1.0	1.0
		5.5	13.0	15.0	8.0	5.0	7.0	5.3	7.0	3.0	7.0	5.0	1.0
		15.0	25.0	30.0	20.8	10.0	13.0	15.0	15.0	9.0	17.5	16.0	10.0
Mixed	<i>1</i>	0.0	1.9	27.1	52.8	0.0	7.0	28.8	50.0	0.0	1.0	15.0	49.2
		1.1	30.0	50.5	71.5	0.1	25.0	45.0	61.0	0.0	15.0	45.0	75.5
		20.0	60.2	77.5	100.0	10.0	30.0	61.1	87.3	0.1	33.0	60.0	99.0
	<i>2</i>	0.0	6.9	20.0	35.5	0.0	7.9	18.0	24.5	0.0	14.5	24.5	33.0
		1.0	18.7	38.5	50.0	0.2	20.0	34.0	40.2	0.1	30.0	40.5	60.0
		10.0	38.0	54.5	77.5	10.0	39.4	50.0	53.5	19.5	42.2	50.2	70.0
	<i>3</i>	0.0	5.0	8.0	8.2	0.0	7.6	10.8	10.0	0.0	9.4	18.5	20.0
		2.0	10.0	18.9	25.0	0.5	17.0	20.0	25.6	5.0	21.0	30.5	40.0
		10.0	35.0	35.0	54.0	9.8	31.4	34.5	35.1	16.0	29.5	44.0	60.0
	<i>4</i>	0.1	2.0	2.0	1.0	0.2	5.0	4.0	1.0	0.1	5.0	5.0	1.0
		2.0	5.2	8.0	10.0	5.0	10.0	7.5	5.0	5.0	11.6	10.0	5.0
		15.0	25.0	15.0	17.8	15.0	23.0	20.0	11.0	15.0	21.6	27.4	20.0
Pooled	<i>1</i>	0.0	1.0	31.0	50.8	0.0	5.0	16.8	40.0	0.0	1.0	20.0	49.0
		1.0	28.0	60.0	77.6	1.1	20.6	33.0	56.4	0.0	15.0	45.0	68.5
		20.0	60.0	90.0	113.2	11.2	30.0	50.2	74.2	0.1	45.3	63.0	97.8
	<i>2</i>	0.0	10.0	20.8	36.0	0.0	7.0	15.0	19.5	0.0	10.0	18.0	30.5
		2.0	25.0	45.0	67.0	2.0	15.5	25.2	34.5	0.0	25.0	35.5	56.0
		15.0	51.0	64.2	85.5	10.0	34.5	45.1	50.0	11.0	40.8	52.0	70.0
	<i>3</i>	0.0	6.7	13.0	12.0	0.0	7.0	10.2	10.0	0.0	8.0	14.0	15.0
		4.0	20.0	27.5	33.5	1.0	12.1	17.6	20.5	4.0	19.0	25.0	30.5
		15.0	38.5	50.0	63.8	10.0	25.8	31.4	35.0	15.0	28.4	40.0	57.5
	<i>4</i>	0.1	4.0	4.0	1.0	0.2	5.0	2.0	1.0	0.0	4.2	1.0	1.0
		3.8	10.0	10.0	9.5	5.0	10.0	7.1	5.0	5.0	10.0	8.0	2.0
		15.0	25.0	21.0	20.0	11.0	17.5	16.0	15.0	11.0	20.0	21.0	15.0

4.2. Fight or Surrender

In analyzing bidding behavior, it is pertinent to identify whether the frontrunner or the underdog is bidding more aggressively at each stage of the tournament—or even if there is a difference. If the underdog is engaging in last stand behavior, then we would expect his bids to increase relative to his opponent’s as he nears an overall loss. We therefore want to compare the underdog’s bid at state (i, j) with the frontrunner’s bid at the symmetric state (j, i) . However, we also want to account for the fact that some players are inherently more aggressive or passive and will thus reach certain states with higher probabilities. To do this, we use a fixed effects regression model with cluster robust standard errors (at the session level) where the dependent variable is a player’s bid.

To account for the different states within the tournament, we include a set of dichotomous variables, equal to one if a bid is made from that particular state and zero otherwise. Since we are interested in comparing bidding behavior at state (i, j) with state (j, i) , we take advantage of the fact that when a categorical variable with N distinct values is represented by a set of $N - 1$ dichotomous variables, the coefficients of the $N - 1$ dichotomous variables can be interpreted directly in reference to the omitted N^{th} value. Thus, we are interested in the coefficient for state (i, j) in a regression where (j, i) is the omitted state.²² Letting $\mathbf{s}(\mathbf{j}, \mathbf{i})$ be the vector of dichotomous state variables which omits state (j, i) , we use the following fixed effects model to predict player k ’s bid at time t within the experiment:²³

Model 1.
$$\widehat{bid}_{k,t} = \beta_0 + \mathbf{s}(\mathbf{j}, \mathbf{i})'_{k,t} \boldsymbol{\beta}_s + f_k + \varepsilon_{k,t}$$

Table 6 shows the Model 1 coefficients for bids made at (i, j) relative to (j, i) where $i > j$. A common result that holds throughout most of the treatments is that players tend to bid more aggressively when they are behind, and they bid increasingly more aggressively as they fall farther and farther behind. For instance, in the mixed group of the Win 15 Lose 285 treatment, players tend to bid 5.61 more rupees at $(4, 3)$ than they would at $(3, 4)$; this difference then increases until by $(4, 1)$, players are submitting bids that average 15.43 rupees higher than at $(1, 4)$. Similar increases can be seen in each treatment of the experienced group, although the magnitudes are not quite as large. Given our last stand hypothesis (H1), we would expect to see this type of behavior in the Win 15 Lose 285 treatment and to a lesser degree

²²Alternatively, for two states that are not omitted, we could obtain the relative difference by subtracting one of the coefficients from the other, taking care to compute the appropriate standard error for the difference.

²³There are two time components: the tournament number and the bids within each tournament. Since the number of bids per tournament may vary between one and seven, we interpolate the timing of each bid to be at one-seventh intervals between tournaments.

Table 6: Bidding Behavior at State (i, j) Compared to (j, i) : Model 1 (baseline)

		Win 15 Lose 285			Win 150 Lose 150			Win 300 Lose 0		
		4	3	2	4	3	2	4	3	2
Experienced	1	6.37 (3.00) > 0**	-4.27 (6.35)	-0.95 (3.77)	5.63 (2.15) > 0***	4.01 (1.97) > 0**	-1.26 (4.11)	10.31 (2.81) > 0***	-0.65 (2.17)	-0.53 (1.69)
	2	4.56 (1.73) > 0***	4.47 (0.99) > 0***		4.92 (1.79) > 0***	0.90 (1.37)		6.28 (0.83) > 0***	5.12 (2.88) > 0**	
	3	1.32 (2.36)			2.02 (2.11)			2.67 (0.13) > 0***		
Mixed	1	15.43 (1.15) > 0***	14.23 (4.14) > 0***	4.10 (4.15)	4.31 (5.16)	3.47 (4.36)	5.38 (2.74) > 0**	1.73 (2.00)	-10.14 (3.41) < 0***	-13.99 (7.71) < 0**
	2	9.62 (1.84) > 0***	5.58 (0.87) > 0***		0.66 (4.03)	0.55 (2.21)		5.01 (0.96) > 0***	-1.19 (7.40)	
	3	5.61 (0.52) > 0***			-3.80 (1.80) < 0**			2.55 (1.90) > 0*		
Pooled	1	10.91 (2.43) > 0***	4.32 (5.64)	1.29 (2.60)	5.06 (2.79) > 0**	3.82 (2.19) > 0**	2.24 (2.57)	6.33 (2.53) > 0***	-5.08 (2.66) < 0**	-6.69 (4.20) < 0*
	2	6.72 (1.55) > 0***	4.78 (0.61) > 0***		2.69 (2.33)	0.74 (1.11)		5.73 (0.63) > 0***	1.94 (3.76)	
	3	3.19 (1.55) > 0**			-0.82 (1.85)			2.67 (0.80) > 0***		

Significance levels: * 10%, ** 5%, *** 1%.

in the Win 150 Lose 150 treatment. It is somewhat surprising, however, that the regression for the Win 300 Lose 0 treatment produces similar results. In both the mixed and experienced groups of the Win 300 Lose 0 treatment, players submit bids at (4, 3) and (4, 2) that are significantly greater than the bids they would place at (3, 4) and (2, 4). In the experienced group, the coefficient for bids at (4, 1) relative to (1, 4) is a shockingly large and significantly positive 10.31. Given that the overwhelming majority of bids at (4, 1) are at or near zero, and that the median bid at (1, 4) is one, this coefficient is largely driven by outliers. The mixed treatment, however, does appear to exhibit some of the surrendering behavior we would expect to see when losing is costless (H2). At (3, 1) and (2, 1) players bid 10 to 14 rupees less than they would at (1, 3) and (1, 2).

There is little if any evidence from these regressions of players defending their lead in the tournament (H3). Theoretically, this should happen when players are ahead by one state. This may explain the significantly negative coefficient at (4, 3) in the mixed Win 150 Lose 150 treatment, and it is also possible that the surrendering behavior at (3, 1) and (2, 1) in the mixed Win 300 Lose 0 treatment is in fact a defense of the lead. Finally, it is worth noting that the regression coefficients for the pooled data are roughly an average of when the regressions are run separately and that statistical significance carries over in most cases.

4.3. Winning Margins and Initial Leads

Our final two hypotheses address the implications of making a last stand versus surrendering—specifically in terms of the size of the winning margin (H4), and also in terms of the importance of winning the initial battle of the tournament (H5). A natural consequence of last stand behavior is that the size of the winning margin tends to decrease. Landslide victories, on the other hand, frequently occur when players are prone to surrendering. Table 7 reports the distribution of winning margins for completed tournaments within each treatment.²⁴ The most pronounced differences in these distributions occur at the endpoints where the winning margin is either one or four battles. In support of our hypothesis, landslide victories are clearly more pronounced in the Win 300 Lose 0 treatment. There, landslide victories account for over 40% of completed tournaments, while the number is between 27% and 31% in the treatments with losing penalties. Neck-and-neck victories likewise increase substantially from the Win 300 Lose 0 treatment to the Win 15 Lose 285 treatment—increasing by 5.5 percentage points in the experienced group and 12.3

²⁴The counts underlying the percentages in Table 7 have been adjusted to account for attrition with the random ending rule. That is, the observations for a winning margin of k have been multiplied by $1/\delta^{4-k}$.

Table 7: Winning Margins by Treatment (in %)

		Winning Margin			
		4	3	2	1
Experienced	L285	27.2	24.1	27.6	21.1
	L150	27.0	27.3	29.4	16.3
	L0	41.0	22.8	20.6	15.6
Mixed	L285	29.3	26.0	22.6	22.1
	L150	31.4	27.4	16.5	24.7
	L0	40.5	28.2	21.5	9.8
Pooled	L285	28.2	25.0	25.2	21.6
	L150	29.4	27.4	22.4	20.9
	L0	40.7	25.4	21.1	12.8

percentage points in the mixed group.

The initial battle is often viewed as pivotal in deciding the ultimate outcome of a dynamic contest.²⁵ It is advantageous in Gelder (2014) but decisive in Konrad and Kovenock (2009). In Gelder’s framework, the probability of an upset is increasing in the relative size of the losing penalty (as is the strength of the last stand). To a degree this is apparent in the data as well. The first thing that Table 8 illustrates is that winning the initial contest is a strong correlate of winning the ultimate tournament—or at least being in the lead at the time the tournament ends (whether by winning or through the random ending rule). Across the different treatments, roughly 70% to 80% of all winners at state (4, 4) went on to win the tournament. Second, the Win 15 Lose 285 treatment appears to have the highest degree of upsets. This is clear in the experienced group where the Win 15 Lose 285 treatment is several percentage points below the other two treatments. The mixed group is less clear since the Win 15 Lose 285 treatment is only lower than the other two treatments when tournaments that ended early are factored in.

²⁵For example, Klumpp and Polborn (2006) examine the disproportionately large amount of attention that New Hampshire and Iowa receive as the first states to vote in the US presidential primary elections.

Table 8: Percent of Initial Battle Winners to Win the Tournament

		Win Tournament	End in Lead
Experienced	L285	68.5	70.6
	L150	78.2	78.3
	L0	76.3	77.2
Mixed	L285	79.0	75.0
	L150	79.4	79.4
	L0	78.9	79.4
Pooled	L285	73.7	72.8
	L150	78.8	78.9
	L0	77.6	78.3

5. Individual Behavior and Player Types

The analysis thus far has taken an aggregative view of the experimental data. We find, however, that many of the anomalies, as well as expected results at the aggregate level, are illuminated by a careful study of individual player behavior. Although we have insufficient information to determine the behavior strategies employed by players in each tournament, we are able to examine each player’s *realized strategy*—that is, the set of history-contingent actions taken by the player along the realized path of play 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 will often conform to the same category repeatedly. Furthermore, the same categories of realized strategies independently arise in the different experimental sessions. The heterogeneity of realized strategies and the implications for the response of one player’s behavior strategy to another’s are key to understanding behavior in this experiment.

As an initial example, consider Figure 4 which illustrates bidding behavior along the exterior states of Figure 1 where one player has either consistently won or lost.²⁶ Several of the pervasive categories of realized strategies and their implied bidding patterns can be identified in this figure. The top row plots cases where a losing trajectory is followed and (4,1) is reached. Bids at (4,3) are shown on the abscissa

²⁶A few outlying points have been omitted to focus these graphs on the area of greatest concentration. The points have been slightly perturbed to reduce overlap.

and the subsequent bids at (4,1) on the ordinate. The points are further coded (see the key to the right of the figure) by whether the player's bid at (4,2) increased, decreased, or remained equal to the bid at (4,3). The bottom row of the figure displays cases where a winning trajectory is followed and (1,4) is reached, similarly plotting bids at (3,4) against those at (1,4).

A striking feature of the bids along the losing trajectory is the concentration of points along the axes and at the origin, with relatively few bids in the interior of the graph. There are four distinct categories of realized strategies represented here. First, points on or near the y-axis represent a last-stand strategy entailing a token bid at (4,3), frequently another at (4,2) (as indicated by the prevalence of black triangles—see the key to the right), and then a sharp increase at (4,1) in the face of a tournament loss. Note that the last stand strategy is much more prevalent in the treatments with a losing penalty than in the Win 300 Lose 0 treatment. The second realized strategy is one of becoming passive and ultimately giving up as resistance mounts. It is represented by the points along the x-axis and away from the origin, indicating that these players were actively competing at (4,3). A few of these players rallied at (4,2) (as shown by the red squares), while several others began to give way (the blue circles). By (4,1), however, they had all given up. A third realized strategy, represented by the cluster of points at the origin, is to give up from the beginning. This is the maximin strategy in the sense that, while anticipating a tournament loss, players maximize their payoff by bidding zero. The fourth strategy—escalation—is represented by the red squares in the interior of the graphs. These are players who were actively competing and increasing their bids, yet continued to be outbid.

Bidding patterns along the winning trajectory clearly interact with those along the losing trajectory. Players either escalate, increasing their bids at (2,4) and (1,4) relative to (3,4); or they curtail their bids as they move toward victory. Since players in the second case continue to win even as they reduce their bids, they are clearly adapting to the strategy of a less aggressive rival. Although some rivals remain passive to the end, the last stand strategy frequently takes advantage of players who have been lulled into thinking they can win at (1,4) with a minimal bid. Escalation, the other major behavior along the winning trajectory, likely occurs either as a preemptive move or when faced with an aggressive competitor.

In order to further explore the realized strategies observed throughout the experiment, and not just those along the two specific trajectories in Figure 4, we individually looked at each of the last ten tournaments for all 216 subjects in the experiment. Examples of complete tournament data for six subjects are shown in

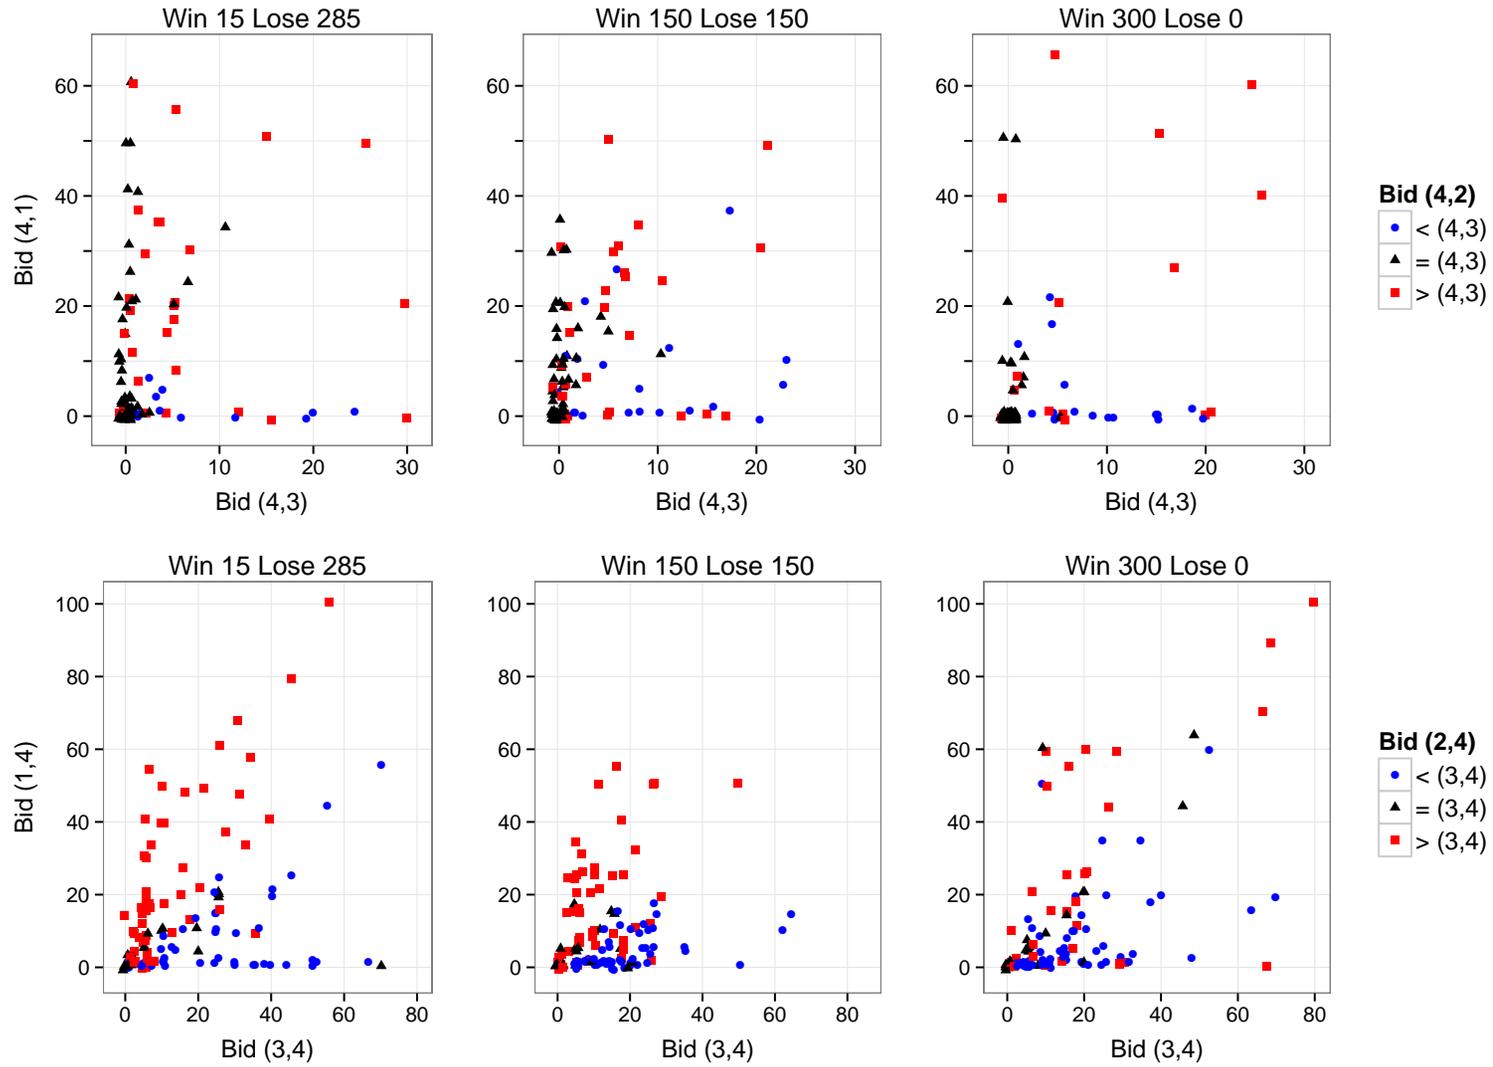


Figure 4: Bids along the trajectories leading to (4,1) (top row) and (1,4) (bottom row)

Tables 9 through 14 (these tables will be described in detail later). Similar to the results from Figure 4, we identified four major realized strategies that arose frequently and repeatedly across the different experimental sessions: maximin, last stand, escalate when challenged, and passive when challenged. There were, however, nuances to these strategies, the most prominent of which was the level of a player’s bid at the initial state (4,4). The initial bid served as a signal to some degree and most players varied it only slightly from tournament to tournament.²⁷ A high initial bid typically signaled a fairly aggressive player; low initial bids, on the other hand, carried less information as they were commonly used by both aggressive and non-aggressive players. To capture this feature, we divided players not only by the category of realized strategy they predominantly used, but also by whether their initial bids tended to be low, moderate, or high.²⁸

From a strictly game theoretic standpoint, we should note that our classification of player types by realized strategies is necessarily limited, since it is based only on behavior in subgames that actually appear in the data. Furthermore, we are not able to distinguish between an action chosen deterministically or through randomization. We do not observe the frequency of actions taken when a subject is confronted with every conceivable history in every stage of the game. Rather, we observe how players actually bid as they face opponent after opponent in one tournament after another and with varying histories in each tournament. Given the rich amount of variation that arises with 216 subjects in 18 experimental sessions, there are some unmistakably consistent patterns of player behavior—consistent enough that we believe the classification of player types is justifiable.

Developing a taxonomy of the different realized strategies was an iterative process. We began by identifying the four recurring qualitative categories in a sample of the data by poring over tables, such as those illustrated in Tables 9 through 14, showing the bid of a given subject and his rival’s in each realized state of each tournament in which the subject played. This process led to a loose definition of the recurring strategies. Working independently, we then classified each player by their most commonly used strategy (if that was at all clear). Differences in our independent classifications led to more rigorous definitions of the strategies. The final taxonomy,

²⁷For 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 were matched at random for each tournament.

²⁸We defined 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 fairly common. For instance, 0, 1, 2, 5, and 10 account for 54% of all initial bids, with 0 alone counting for 23%.

presented here, is specific enough that a computer can apply these rules to classify the category of realized strategy used by a player in each tournament.

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 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 $(4, 1)$.*²⁹
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 in any battle place a bid 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 satisfied along the realized path of the tournament: The player reaches some state in which he is both strictly behind in the tournament and has two or more losses. In that state, the player places a bid that is 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.

²⁹These exclusions are meant to distinguish escalation in the face of a challenge from last stand behavior.

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.³⁰ Of the 216 players, 201 can be classified by their initial bid, 180 by a realized strategy, and 169 by both.³¹ Examples of players who are classified by each of the four categories of realized strategies are shown in Tables 9 through 14. These tables show bids at each state and the resulting tournament paths for the last ten tournaments of play. To the left of each tournament grid is the number (11–20) and outcome of the tournament (L: loss; W: win; E: end early). Within the cells of the tournament grid, a player’s own bid is written on the left and the rival’s on the right. Each of the players represented in these tables has been selected from a different session.³²

Table 9 is a classic maximin player who always bids zero, regardless of how low his opponent might bid. Table 10 shows a player who is passive when challenged, a behavior evident in tournaments 12, 16, 18, and 19. This table also illustrates the significance of the term *when challenged*. There is no opposition in tournaments 11, 13, and 17, so the player resorts to placing minimal bids—a widespread behavior that does not delineate a unique enough aspect of the player’s realized strategy to be classified. Tournament 20 likewise cannot be classified because of its early

³⁰A tournament needed to continue for a minimum of three battles before any classification could be made, and 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.

³¹There 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 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. For example, one player had five tournaments classified under the last stand realized strategy and five classified under escalate when challenged. Eight players had three or more tournaments classified under each of two different realized strategies; nine players had two or more tournaments apiece under three different realized strategies; and four players had at least one tournament classified under each of the four different realized strategies.

³²Sessions are labeled by their losing penalty (L285, L150, L0) and subject pool (E: experience; M: mixed), as well as by a session letter (A, B, C). For instance, session C of the Win 150 Lose 150 mixed treatment is denoted L150 MC.

Table 9: Maximin Player from Session L150 MC

11	0 8					12	0 10				
	0 7						0 10				
	0 6						0 25				
L	0 6					L	0 25				
13	0 1					14	0 1				
	0 1						0 5				
	0 5						0 10				
L	0 20					L	0 1				
15	0 0.1					16	0 1				
	0 5						0 1				
	0 15						0 10				
L	0 30					L	0 1				
17	0 0.1					18		0 0	0 0		
	0 0.1						0 0	0 1			
	0 0.1						0 1				
L	0 0.8					L	0 0				
19	0 0.3					20	0 0.9				
	0 11						0 0.6				
	0 0.2						0 0.1				
L	0 9					L	0 10				

Table 10: Passive when Challenged Player from Session L0 EA

11						12		0 10			
								0 15			
								0 19			
W	3 0	4 0	1 0	1 0		L	3 2	5 9			
13						14					
										30 11	
W	4 0	5 0	1 0	1 0		W	4 1	10 1	2 1	2 5	
15						16	0 1				
							0 5				
							0 9				
W	4 1	6 1	2 1	10 10		L	4 5				
17						18	0 1				
							0 3				
							0 11				
W	4 0	5 0	2 0	1 0		L	4 5				
19		0 5				20					
		0 25									
		0 31									
L	4 2	10 12				E	4 0	5 0			

Table 11: Last Stand Player from Session L150 MB

11							
E	0	8					
13	20	2	20	31			
L	0	15					
	0	10					
L	0	14					
15	20	0.2	30	33			
L	0	11					
	0	22					
L	0	16					
17	15	5	30	30			
L	0	5					
	0	5					
L	0	0					
19	30	5	40	38	65	52	90 78
W	0	10					
	0	24					
W	0	16					
12							
E	0	11					
E	0	11					
14	20	15.2	27	75			
L	0	15					
	0	15					
L	0	15					
16	30	15	30	75			
L	0	15					
	0	15					
L	0	15					
18							
E	10	100					
20	20	2	25	66			
L	0	2					
	0	15					
L	0	15					

Table 12: Last Stand Player from Session L285 MC

11	41	2	61	51	100	86	
E	1	4					
	1	2					
E	0	5					
13	1	9					
L	1	3					
	1	24					
L	0	25					
15							
E	0	2					
	0	15					
E	0	25					
17	3	1	45	5	0	7	
L	0	1					
	0	1					
L	0	44					
19	1	41					
L	1	11					
	0	6					
L	0	0					
12	41	2	43	50			
L	1	11					
	1	25					
L	0	19					
14	21	2	61	45	101	100	41 126
L	1	15					
	1	2					
L	0	25					
16	6	5					
E	0	5					
	0	5					
E	0	5					
18			61	101			
L			21	26			
			0	9			
L	0	0	0	2			
20	21	11	100	51			
E	0	20					
	0	36					
E	0	26					

Table 13: Escalate when Challenged Player from Session L150 EC

11							
		50 25	50 50	64 51			
W	20 20	21 30					
13							
E	12 6.1	10 0.1	2 6.1				
15							
			12 10	16 15			
			6 8				
W	12 0	2 0	1 5				
17					48 39		
					24 28		
		8 7	12 9.9	14 18			
W	12 0	4 5					
19							
					14 15		
E	12 11	18 0	5 0	2 10			
12						22 20.1	
					12 4.9	14 15	
W	20 5.1	10 0	2 10				
14							
					10 5.1	13 14.8	
E	12 5.1	3 0	2 5				
16							
					22 20		
					6 5	11 15	
W	12 5	10 0	2 10				
18							
		42 0	16 0	24 0			
		19 37					
W	12 11	13 26					
20					34 0	40 40	
					16 20		
					8 10		
L	12 1	6 0	2 5				

Table 14: Escalate when Challenged Player from Session L0 MA

11							
				70 31	100 49		
	40 9	41 15	25 59				
W	25 33						
13							
W	25 0	10 0	1 0	1 0			
15							
			60 55	75 80			
			40 45				
		25 19	26 36				
L	20 2	10 19					
17							
			70 69.8	80 95.9			
		40 36.4	45 52				
	22 21.7	24 32.8					
L	10 15						
19							
			66 70				
			45 53				
		30 23	35 39				
L	10 2	10 23					
12							
		90 80	95 80				
		60 75					
	41 35	45 50					
E	20 40						
14							
						50 21	
					15 11	15 21	
W	20 5	6 5	6 10				
16							
					53 65		
					45 49		
		23 23	31 35				
L	20 1	10 16					
18							
		25 18.6	30 27.7	35 10.5			
W	10 5.1	10 17.8					
20							
						100 101.4	
					70 56.6	75 85	
	25 24.8	36 35	46 48.9				
L	15 18						

termination. The remaining two tournaments, 14 and 15, meet the qualification of escalating when challenged. Since the four passive when challenged tournaments form a strict majority of the six classifiable tournaments, this player is designated as passive when challenged.

In Tables 11 and 12, players repeatedly use the last stand realized strategy. Although not required by definition, these two players nearly always begin with a bid of zero. The player in Table 11 continues to bid zero, mimicking maximin play, but suddenly springs into action at $(4, 1)$. Evidently expecting this maximin play to continue, the rivals in tournaments 13, 15, 19, and 20 dramatically reduced their bids. The opponent in tournament 14 correctly anticipated the possibility of a last stand and submitted a slightly higher bid. Having bid 15 the previous three times, this opponent may have wanted to avoid being outdone by a bid of 15.1 and so bid 15.2. Even if a last stand is successful at $(4, 1)$ it tends to be dramatically harder to secure a victory at $(3, 1)$ since the rivals have been stirred to action. The last stands in Table 12 are quite similar, although it is worth noting that tournaments 13 and 19 do not meet the definition of a last stand but are rather maximin. A last stand bid must have a credible chance of winning given the previous bids in the tournament.³³

The escalate when challenged realized strategy is demonstrated in Tables 13 and 14. In eight of the tournaments, the player in Table 13 opens with a bid of 12—a bid that is strong enough to consistently win at $(4, 4)$ each time it is used. Gaining the lead from the start, the player submits successively lower bids until a rival poses a threat. Typically this occurs when the player loses a round, but in tournaments 11, 18, and 19, the rival's bid at $(4, 4)$ is close enough that it registers as a threat. Once challenged, the player escalates. However, in escalating, this player is more judicious than the player in Table 14. The difference comes in the rate of escalation. The Table 13 player may have two or three rounds of escalation before hitting a bid of twenty, while the Table 14 player can easily exceed a bid of twenty in the first bid of a five or six round escalation. Ultimately, the Table 14 player expends more in bids than the value of the prize-penalty spread in three of the tournaments (11, 12, and 20) and comes close in another two (15 and 17). Part of the issue is that high initial bids leave little room to backtrack in a bidding war. Another part is knowing when to quit. Even if this player would have won the three tournaments that went to $(1, 1)$, consistently bidding zero would still have yielded a higher average payoff.

³³For tournament 13 to meet the definition of a last stand, the player would have needed to bid at least 12.6 at $(4, 2)$ or at least 2.1 at $(4, 1)$. There is no clear realized strategy in tournament 15 since both last stand and maximin remain possibilities since the tournament ended early.

The distribution of player types within each session is summarized in Table 15. Here we classify players both by initial bid and realized strategy. Several notable patterns are present in this table. One that is immediately apparent is the widespread distribution of player types across each session. At least two distinct realized strategies are represented in every session, with escalate when challenged being the most common. Across the low, middle, or high initial bid ranges, escalate when challenged accounts for at least six of the twelve players in all but two sessions. The remaining three realized strategies jointly comprise about a quarter of the classifiable players and a fifth of all players. The incidence of these remaining players across treatments is particularly telling, however, in terms of the last stand and surrendering hypotheses.

There is a relative abundance of maximin and passive when challenged players in the Win 300 Lose 0 treatment, a fact which helps to substantiate our hypothesis of players surrendering when losing is costless. While the Win 15 Lose 285 and Win 150 Lose 150 treatments have one passive when challenged and three maximin players apiece, the Win 300 Lose 0 treatment has three passive when challenged and ten maximin players. The opposite holds true for last stand players. There are only four last stand players in the Win 300 Lose 0 treatment but nine and eleven respectively in the Win 150 Lose 150 and Win 15 Lose 285 treatments.³⁴

Just as players are classified in Table 15 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 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.³⁵ This practice of shopping around for different strategies can be seen in Figure 5, which contains the cumulative distribution of the number of times players use a given realized strategy during the last ten tournaments. Regardless of treatment, more than one-half of all players have experience using the last stand realized strategy. Usage rates for maximin and passive when challenged sit at around a quarter, and nearly everyone has tried escalate when challenged. Notwithstanding this shopping around, the usage rates of the different realized strategies are fairly delineated across treatments, and the distributions can frequently be ranked in terms of stochastic dominance.

³⁴Including three last stand players that cannot be classified by initial bid, there are five last stand players in the Win 300 Lose 0 treatment and eleven in each treatment with a losing penalty.

³⁵Of the 216 players, 65 used only one realized strategy in the tournaments that can be classified; 101 used two, 45 used three, and 5 subjects managed to use each of the four realized strategies.

Table 15: Distribution of Player Types by Session

		Low			Mid			High			Other
		MM	LS	Esc	LS	Esc	Pass	LS	Esc	Pass	
L285	Mix A	1	2	6		2		1			
	Mix B		2	4		3		1			2
	Mix C		2		2			7	1		
	Exp A	1		2		6		1			2
	Exp B			1	1				5		5
	Exp C	1	1	3		1		1	2		3
L150	Mix A		2	1		6		2			1
	Mix B		1			2		5			4
	Mix C	2		1		1		2			6
	Exp A	1	2	1		5		1			2
	Exp B		2	4		2	1		2		1
	Exp C		2	1		5			1		3
L0	Mix A	1				4		3	1		3
	Mix B	3				6					3
	Mix C	1	1	3		1		3			3
	Exp A	2	1			5	2				2
	Exp B	2	2	1		4			1		2
	Exp C	1		1		1			4		5
Total	L285	3	7	16	3	12		1	17	1	12
	L150	3	9	8		21	1		13		17
	L0	10	4	5		21	2		11	1	18

11 players can be classified by realized strategy but not by initial bid. Escalate: L285MB, L285EA, L285EB, L150MB ($\times 2$), L150MC, L0EC ($\times 2$); Last Stand: L150MB, L150MC, L0MC.

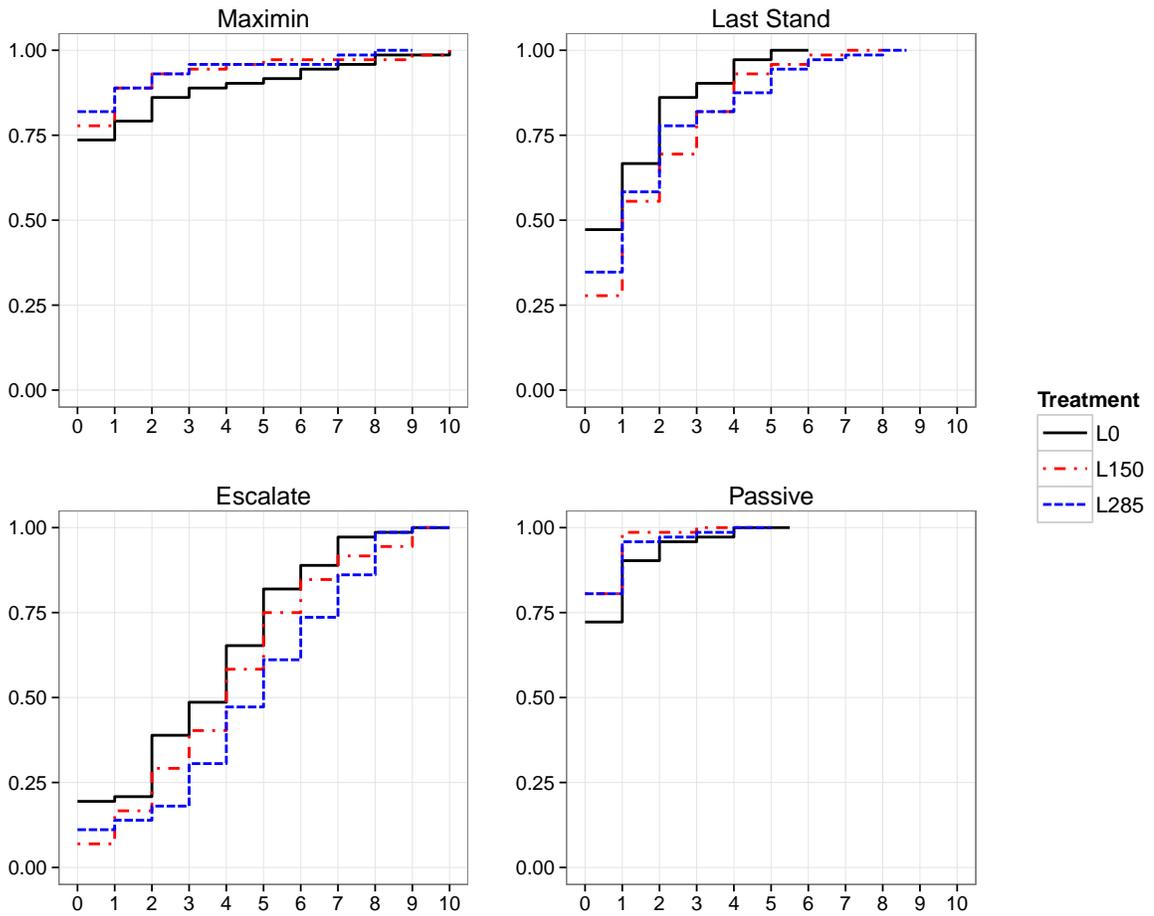


Figure 5: CDF of number of times players use a strategy during last ten tournaments

For maximin and passive when challenged, the Win 300 Lose 0 treatment comes close to first-order stochastically dominating the other treatments. Although the ranking is less clear near the top of the distributions (where the Win 150 Lose 150 treatment actually pulls ahead of the Win 300 Lose 0 treatment for playing maximin eight times), the Win 300 Lose 0 treatment has a pronounced lead earlier on. In the two treatments with losing penalties, only a few players use maximin more than two times, and hardly any use passive when challenged more than once. This seems to indicate that players in the Win 15 Lose 285 and Win 150 Lose 150 treatments are simply trying out these realized strategies rather than committing to them. Even in the Win 300 Lose 0 treatment, players never have more than four tournaments that can be classified as passive when challenged. This lack of observations is likely due to two definitional factors which are out of the player's control: first, that his rival is aggressive; and second, that the tournament continues long enough for him 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. The lead belongs to the Win 15 Lose 285 treatment thereafter. It is fitting that the last stand strategy is employed with greater frequency in the treatment with the larger penalty for losing.

Use of the escalate when challenged strategy is crisply ordered across the three treatments over nearly the entire distribution. The Win 15 Lose 285 treatment has the highest occurrence at nearly every level, followed by the Win 150 Lose 150 treatment, and then the Win 300 Lose 0 treatment. It is not entirely clear why such a clear ordering should exist, although there are several potential explanations. First, even though the prize-penalty spread remains the same across treatments, it may be that players feel that there is more at stake when there is more to lose (as is the case in prospect theory) and therefore behave more aggressively. This may lead to more escalation.

Another possible explanation is that the reduced incidence of escalation in the Win 300 Lose 0 treatment may simply be an artifact of potential observations being censored due to the definition of the escalate when challenged type. To meet the definition, a player must bid weakly greater than his opponent's last bid and strictly greater than his own last bid in two separate battles. Consequently, if a player employs a behavioral strategy of responding to any challenge with an increased bid, but otherwise maintains or lowers his bid, then if the player does not face any challenges from his opponent, the player's realized strategy for that tournament will not be categorized under escalation. Since the incidence of maximin and passive when challenged are highest in the Win 300 Lose 0 treatment, it is likely that such censoring is more problematic because there is a higher probability of facing a player who would not challenge. Although this explanation seems reasonable, it does not appear to account for the ranking of treatments within the data.³⁶

³⁶As a conservative test, we constructed a distribution similar to those in Figure 5 for the number of times a player's realized strategy was either classified as escalate when challenged or remained unclassified in a tournament that did not end after the first or second battle (no tournaments within our taxonomy can be classified prior to the third battle). In essence, by grouping the unclassified tournaments with the escalation types, we are treating the unclassified tournaments as if they were escalation types. We find that there is still a distinct ranking between the Win 15 Lose 285 and the

We now turn to the question of how well each of these strategies performs. Figure 6 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 the escalate when challenged players are also coded by initial bid. Each row denotes a separate session of the experiment, and the sessions are sorted from top to bottom based on the median payoff of the session. The maximin players serve as a baseline for comparison since every player 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. Depending on the actual realization of the number of tournaments that ended early, the maximin payoffs tend to cluster around these points.³⁷

More than any other realized strategy, escalate when challenged with a high initial bid frequently fares worse than maximin. This was the case with the player in Table 14 who received the lowest average payoff in the Win 300 Lose 0 treatment. Sessions where multiple players adopted this strategy had an especially hard time since it led to routine high-stakes bidding wars. These sessions have the lowest median payoffs within each treatment.

It is conceivable that the poor performance of escalating player types with high initial bids might result from the same censoring problem noted earlier. A player with a high initial bid might immediately deter any potential challenge, precluding the need to raise bids at least twice as required by the definition of *escalate when challenged*. In such a tournament, even though the player employs a behavior strategy that escalates in subgames involving a challenge, his realized strategy does not involve escalation. Moreover, because the player deters his rival, the player is likely to have a high tournament payoff, offsetting potential losses that arise in other

Win 300 Lose 0 treatments. The overall ranking of the Win 150 Lose 150 treatment throughout the distribution is less clear, although it does tend to the middle. Another factor which likely effects the incidence of escalation across treatments in Figure 5 is the realization of the random termination of tournaments. The Win 300 Lose 0 treatment had more than expected, with 148 tournaments ending after the first or second battle; the Win 150 Lose 150 had 132; and the Win 15 Lose 285 had only 108.

³⁷The occasional maximin player would win a tournament—frequently by deviating from the maximin realized strategy when winning looked feasible. There are a few scattered instances where a tournament was actually won with a cumulative bid of zero. Interestingly, there is only one such case where both the winner and the loser were classified as maximin players.

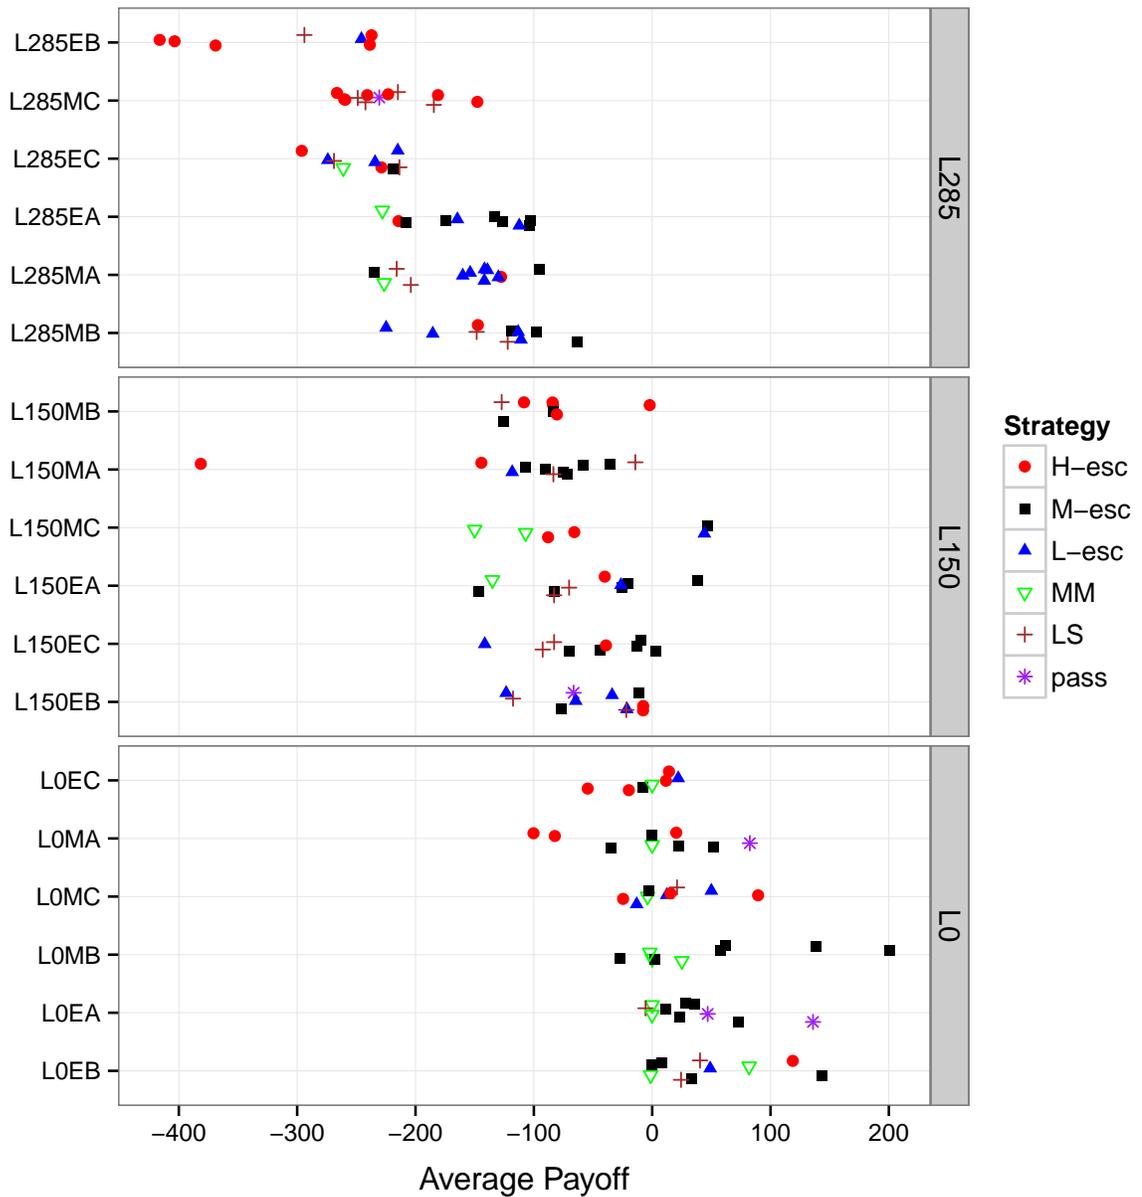


Figure 6: Average payoff by player type and session. Sessions are ranked by median payoff.

tournaments where escalation actually occurs.

It turns out that this effect does not appear to drive the results. First, the poor performance of escalating types with high initial bids appears to be as much a feature of the treatments with a high losing penalty as the treatment with no penalty. Since the probability of facing surrendering behavior is not as large in the high penalty treatments, it does not seem to be a factor in the payoff rankings. More to the point, Figure 6 illustrates player types and not the incidence of individual realized strategies. For a player to be categorized as an escalating type with a high initial bid, he must be categorized as playing an escalating strategy at least three times and in a strict majority of all categorized tournaments. He must also place a high initial bid in a strict majority of the tournaments. An escalating player placing a high initial bid who sometimes succeeds in preempting may still meet this requirement over the course of the last ten tournaments. When he does not, it would likely mean that preemption occurs sufficiently often that the player is not categorized by type and hence appears among players listed as “other” (see Table 15). A careful examination of the realized strategies of all subjects listed as “other” yields little evidence that this censoring is an issue affecting the payoff rankings.

Although escalation can be a harmful strategy 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 for bidding wars, and it also frequently results in an early lead in the tournament. Low-range initial bids, on the other hand, frequently forfeit the early lead. This is likely why players who escalate from a low initial bid have lower payoffs than those who escalate from the mid-range.

Last stand players, overall, do marginally better than the maximin players. Instead of simply swallowing the losing penalty, the last stand strategy seeks to buy the player time so that the tournament may end early. A last stand is profitable relative to maximin if the bidding cost from successfully delaying the tournament is less than the penalty for losing. Due to the lack of observations, it is difficult to assess the passive when challenged realized strategy. However, it appears to do very well in the Win 300 Lose 0 treatment where two of the players using it rank first in their sessions. These players are vested enough in the tournament to win against easy competitors, but they are also quick to avoid costly confrontation.

We also examined performance across all tournaments ranked solely on the classification of the magnitude of a player’s initial bid at (4,4). It turns out that this

ranking is consistent with our ranking of performance by player types above. In the Win 300 Lose 0 treatment players placing an initial bid in the high range earned a mean payoff of 15 rupees, versus a mean payoff of 22 rupees for players initially bidding low, and 34 rupees for those initially bidding in the middle range. Although high initial bidders in the Win 150 Lose 150 treatment performed better than low initial bidders (a mean payoff of -75 rupees versus -81 rupees), they fared worse than players with initial bids in the middle range (a mean payoff of -60 rupees). Finally, high initial bidders fared even worse in the Win 15 Lose 285 treatment, where they earned a mean payoff of -238 , versus -185 for low initial bidders, and -181 for players bidding in the middle range. The main conclusion from this exercise appears to be that, regardless of whether a player's behavior conforms to a realized type or not, placing an initial bid in the middle range yields the highest payoff, while placing an initial bid in the high range is relatively unprofitable.

In ending, there are a few noteworthy connections between demographics and the realized strategies that players adopt. Females have a much higher propensity to be classified as using the escalate when challenged strategy than males (75% of females compared to 49% of males). Combined with initial bidding levels, the gender difference vanishes among low initial bidders but is most pronounced for those placing high initial bids. Thus, here, as in other contest experiments, women tend to bid more aggressively than men (see Dechenaux et al. 2014). Men, on the other hand, were much more likely to use the last stand or maximin strategies. Politically, the usage rate of escalate when challenged peaked on each extreme of the political spectrum and fell to a trough in the middle. The last stand was a more politically centrist strategy. Risk aversion also played a role with more risk averse subjects tending toward escalation slightly more than their counterparts, while the last stand had its highest frequency of use among those with the greatest tolerance for risk.

6. Conclusion

Using a controlled laboratory experiment, we examine how aggressively players compete at different stages of a best-of-seven tournament in which each component contest is an all-pay auction. We find that bidding behavior differs considerably based on whether or not players face a substantial penalty for losing the tournament.

When the cost of losing is high, players who have fallen behind frequently submit fairly large bids in an effort to avoid a tournament loss. This is consistent with the theoretical predictions of Gelder (2014) and embodies the notion of a last stand.

Within our experiment we specifically identify several players whose realized strategies routinely conform to last stand behavior. Moreover, the incidence of last stand behavior is much more prevalent in the treatments with losing penalties. Conversely, we find that subjects are more prone to surrender in tournaments that have no losing penalty. Many subjects even employ a maximin realized strategy of consistently bidding at or near zero—in essence, surrendering from the start. Although forfeiting the tournament prize by refusing to compete, these players also avoid undue expenditure. A second type of surrendering behavior found in the data, albeit with lower incidence, consists of actively competing until a player faces strong resistance, at which point the player surrenders.

The most prevalent player type we identify is to escalate the conflict through a bidding war. In fact, players employing this strategy are among the most successful and unsuccessful players that we observe. Players who are overly aggressive, starting with a high bid and then continuing to escalate, tend to perform the worst in terms of average payoffs, while moderately aggressive players who start at intermediate bids and then escalate often do the best.

Equilibrium behavior dictates that a player will defend his lead in the overall tournament by bidding more aggressively than his opponent whenever his opponent is trailing behind by a single battle victory. We find little evidence to support this hypothesis. However, we do find that the theoretical predictions for winning margins are substantiated. Neck-and-neck tournaments are more likely when the losing penalty is relatively high, and landslide victories occur more frequently when the winning prize is large relative to the losing penalty. Finally, tournament upsets are fairly uncommon, with the winner of the initial battle going on to win the entire tournament roughly three-quarters of the time.

There is a natural challenge in designing an experiment which exposes subjects to a loss that is in some sense meaningful, yet minor enough to comply with standard institutional review board guidelines. We suggest, therefore, that any evidence of last stand behavior in our low stakes experiment would be magnified in situations that involve more substantive losses. It is also instructive that escalation is nearly the default behavior in our experiment. Although this behavior may be socially desirable in settings such as sporting events, it may be highly undesirable in other environments such as legal disputes.

The player types we identify arise through an examination of the realized paths of play across experimental sessions, as opposed to being induced through the presence of a player with a known behavioral strategy (see Embrey et al. 2014). Method-

ologically, the analysis of realized strategies and player types presented here can be extended to a host of other experiments involving dynamic games—especially those with rich action spaces in which it is difficult to ascertain (or even completely describe) a subject’s full extensive form strategy. Studying individual subjects in this way can reveal numerous patterns that are easily lost in the aggregated data. Although it is clearly beneficial to conduct experiments which attempt to solicit a more complete picture of players’ extensive form strategies, we believe that the examination of the types generated by realized play is a useful tool for future experiments.

Appendix A. 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:

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

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.

At the conclusion of the experiment, subjects were given the following demographics survey.

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]

Acknowledgments

We express gratitude to Jeffrey Kirchner for programming the experiment and gratefully acknowledge helpful comments and suggestions from Gabriele Camera, Brice Corgnet, David Porter, Stephen Rassenti, Roman Sheremeta, Vernon Smith, Nathaniel Wilcox, Bart Wilson, and John Wooders. This work has also benefited from participants at the 14th APET Conference (Seattle, Washington); the 25th International Conference on Game Theory (Stony Brook, New York); and seminars at the Economic Science Institute (Orange, California) and the Max Planck Institute for Tax Law and Public Finance (Munich, Germany).

References

- Agastya, M., McAfee R.P., 2006. Continuing wars of attrition. Unpublished manuscript.
- 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., 2014. A survey of experimental research on contests, all-pay auctions and tournaments. *Exper. Econ.*, doi: 10.1007/s10683-014-9421-0.
- 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., 2014. Bargaining and reputation: An experiment on bargaining in the presence of behavioral types. *Rev. Econ. Stud.*, doi: 10.1093/restud/rdu029.
- 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. Unpublished manuscript.
- 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. Unpublished manuscript.
- Klumpp, 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. Unpublished manuscript.
- 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.
- 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.