

How Hard to Fight? Strategic Sophistication in an Asymmetric Contest Experiment

Stephen Chaudoin*
University of Illinois

Jonathan Woon†
University of Pittsburgh

December 28, 2015‡

Abstract

Many political phenomena—from wars to elections and lobbying—involve winner-take-all contests in which the value of the prize differs across the actors involved and from one issue to the next. To better understand competitive behavior in such environments, we conduct a controlled laboratory experiment in which participants face a series of asymmetric prize values in a lottery contest game. We find support for some, but not all, of the game’s comparative static predictions. Most subjects respond to changes in their own values, but few subjects conditionally respond to cross-player changes. Our data therefore suggest a new type of heterogeneity in the degree of strategic sophistication. We also administer two information based treatments, feedback and a calculator, finding that feedback on past play has a stronger effect on increasing subjects’ strategic sophistication than a payoff calculator. (Abstract: 136 words)

JEL Codes: C72, D72, D74

*chaudoin@illinois.edu

†woon@pitt.edu

‡We appreciate helpful feedback from Ryan Brutger, Soenke Ehret, Ju Yeon Park, Mattias Polborn, Stephanie Wang, Alistair Wilson, and audiences at the American Political Science Association Annual Meeting, the North American Economic Science Association, the University of Pittsburgh, and the University of Illinois.

In domestic and international politics, political actors compete with one another to achieve mutually exclusive goals. These situations are often thought of as contests in which the chance of winning the prize is a function of the effort or resources that each side devotes to it, but where increased competition is socially wasteful. War, for example, is a contest in which devoting more manpower, materiel, and industrial capacity increases the chance of winning but costs human lives and wasted economic productivity (e.g., Slantchev, 2010; Besley and Persson, 2011). Less violently, NGOs influence human rights policy by mobilizing and counter-mobilizing for and against reforms (Bob, 2012; Sell and Prakash, 2004). Regulations governing foreign and domestic interactions, such as tariff or environmental policy, can be thought of as the outcome of competition between special interest groups with opposing preferences who make campaign contributions (e.g., Grossman and Helpman, 1994; Goldstein and Martin, 2000). More generally, rent-seeking and lobbying are thought of as classic examples of contests (Tullock, 1967; Krueger, 1974; Becker, 1983). Similarly, the struggle for control of government between incumbents and challengers has been modelled as a contest in both autocracies (Myerson, 2008; Svobik, 2009) and democracies (Iaryczower and Mattozzi, 2013; Meirowitz, 2008; Serra, 2010).

Our study is motivated by two observations about contests occurring in the real world. First, contests are often *asymmetric*. That is, parties to a conflict can differ in the value they place on winning the contest or in the effectiveness with which they transform their resources into competitive advantages. Winning access to a new ocean port is much more valuable to a landlocked country than a coastal one. Likewise, winning influence over tariff policy is much more valuable to an import-competing firm facing bankruptcy from foreign competition than it is to a consumer who dislikes marginally higher prices. Hiring a lobbying firm with extensive networks may be more effective than spending the same amount with a less connected firm. In electoral contests, incumbency can affect a politician's ability to turn campaign resources into electoral success while

campaign spending has greater diminishing marginal returns for non-incumbents than for incumbents.

Second, contests are affected by *shocks*, meaning that features of the competitive environment can vary between contests and across time. The discovery of natural resources in disputed regions has significant implications for the value of controlling a piece of territory (Ross, 2004*a,b*), as with the presence of oil in the conflict between North and South Sudan or the discovery of alluvial diamonds in Sierra Leone. Actions by international institutions can change costs and valuations in a domestic political contest over whether to comply with the rules of an international organization (Chaudoin 2015). Economic shocks, like the Great Recession, affect a firm or special interest group's urgency of obtaining a protective tariff or favorable regulatory ruling (Henn and McDonald, 2014; Davis and Pelc, 2012). Climate change can affect competition for scarce resources by raising the value of those resources (see Salehyan, 2008; Nordås and Gleditsch, 2007). In the electoral arena, an unforeseen scandal or sudden crisis can shift the advantage from one candidate to another (Abramowitz (1991); Welch and Hibbing (1997); Levitt (1994), while a judicial ruling on campaign finance reform can unexpectedly increase the costs of influencing voters' perceptions (Meirowitz, 2008).

We use a laboratory experiment to study how these features of real-world contests—asymmetry in and changes to players' valuations of the prize—affect behavior in a Tullock-style lottery contest. The laboratory setting is appropriate because we can precisely control these valuations and observe participants' effort levels, quantities which would be difficult to measure in observational settings.¹ While there is a large body of experimental work studying contests in the laboratory, we are the first (to our knowledge) to test the effects of asymmetric and changing valuations on behavior.² Specifically, we test three comparative static predictions derived from Nash equilibrium

¹As with any method, there are trade-offs. Lab experiments are similar to models in that the design focuses attention on a small set of key variables—the setting is not meant to perfectly replicate reality or generalize to observational phenomena. Yet, the experiment is valuable because it allows us to explicitly test behavioral predictions in a controlled setting in which we manipulate the salient features of many real-world situations: asymmetries in/shocks to the value of winning the prize.

²For an extensive, recent survey, see: Dechenaux, Kovenock and Sheremeta (2014). Much of this work focuses on

analysis. First, an increase in one player's own prize valuation directly increases their own effort. However, that increase also has second-order cross-player effects, indirectly causing an increase or decrease in her opponent's effort levels, depending on the players' relative valuations. If it increases her opponent's effort, because her opponent seeks to discourage or deter her from further competition, we call this "doing the deterring." If it decreases her opponent's effort, because the contest becomes more lopsided and the marginal return to effort decreases for her opponent, we call this "getting deterred," a concept similar to the discouragement effect documented in existing experimental work (Gill and Prowse, 2012*b*; Deck and Sheremeta, 2012).

Motivated by a common finding in the literature that players exert too much effort relative to Nash predictions, our design also assesses the degree to which information enhances strategic behavior. In the baseline condition, players know only the rules of the game and the value of the prize to each player. In the *Feedback* treatment, players also observe the effort levels of their opponent and the outcome of the contest after each round. In the *Calculator* treatment, players are given a payoff calculator which allows them to search the action space and observe their own and their opponent's expected utilities for pairs of effort levels. Whereas the feedback treatment gives them *empirical* data about how their opponents have played and their resulting payoffs, the calculator treatment allows them access to *hypothetical* data about payoffs, a tool which is potentially very powerful, but which puts the onus on the participant to take advantage of it.

We find strong support for two comparative static predictions, and mixed support for a third. First, increasing a player's valuation increases their own effort under all treatment conditions. Second, we find support for the "getting deterred" effect, whereby increasing player i 's valuation *decreases* player j 's effort in all but two treatment conditions. In contrast, there is more modest support for the "doing the deterring" effect, as players do not always increase their effort levels in response to increases by their opponents.

explaining the total effort levels and heterogeneity between individuals in contests with stable, symmetric valuations. Work on variation across individuals has focused on a wide array of individual-level explanations, such as demographic characteristics, preferences toward risk, or other behavioral phenomena such as "the hot hand" fallacy.

In terms of information effects, we find that feedback and the payoff calculator both lower effort levels, which brings observed effort levels closer to Nash predictions. However, the magnitude and statistical significance of the feedback effect are much higher than that of the calculator's effect. Even accounting for players' learning over time, the feedback treatment significantly decreases the distance between observed behavior and the Nash prediction. Additionally, we find that the feedback treatment, but not the calculator treatment, increases the number of respondents whose behavior comports with the Nash comparative statics.

In our analysis, we also assess the degree of variation in strategic sophistication displayed by the subjects. A large body of literature, such as work based on level-k models of iterated reasoning (Nagel, 1995; Stahl and Wilson, 1995) or models of cognitive hierarchies (Camerer, Ho and Chong, 2004), classifies individuals based on their degree of strategic thinking. We document a form of variation in strategic thinking that diverges from that of existing work. Specifically, we assess whether there is variation in the degree to which individuals behave according to the comparative statics, and we find that there is indeed significant heterogeneity across individuals in their strategic responses to differences in valuations. A small number of subjects display behavior that is consistent with all of the comparative static predictions while approximately half display behavior that is consistent with only one, but not both, of the cross-player hypotheses.

The implications of this research are broad. First, since contests with asymmetry and shocks are ubiquitous, it is inherently important to assess how individuals approach these situations. Understanding how individuals respond in the lab provides a foundation to assess how individuals may approach these situations in the real world. Second, the feedback and calculator findings shed light on how individuals gain knowledge about their strategic environments. Tangible, experiential information, as embodied by the feedback treatment, more effectively induces strategic behavior than the abstract information embodied in the calculator treatment. This suggests that adversaries who engage each other frequently across multiple contests in the real world will behave more strategically than competitors in one-off interactions.

Third, our documentation of a distinct type of heterogeneity in individual behavior is important since a growing body of literature is interested in an individual's level of strategic sophistication as an explanatory and as an outcome variable. Iterated reasoning and the ability to anticipate your opponent's moves are only two aspects of "strategic thinking." Our research highlights how individuals vary in their ability to understand how changes to the game affect themselves and their opponents incentives. This is a type of dynamic strategic thinking, where individuals vary in how their degree of understanding when it comes to outside shocks that influence strategic interactions. This dynamic aspect of strategic thinking may be more appropriate in real world settings where the parameters of the situation are fluid and strategic reactions are paramount.

Contest Model

We consider a simple contest model in which two players can each exert costly effort in order to increase their chances of winning a prize. Each of the two players, i and j , has a strictly positive value to winning the prize, V_i and V_j , where the prize values are distinct, $V_i \neq V_j$. Their effort levels are denoted e_i and e_j and the players have constant marginal costs of effort, $c_i = c_j = 1$. The contest is a function which maps their effort levels into the probability of winning the prize. The probability that player i wins the contest is $\phi_i(e_i, e_j) = \frac{e_i}{e_i + e_j}$, and we assume that $\phi_i(0, 0) = 0$. This is the familiar ratio or Tullock (1967) contestation function. Player i 's objective function is thus:

$$\Pi(e_i, e_j) = \phi_i(e_i, e_j)V_i - e_i.$$

The Nash equilibrium effort level obtained from the players' accompanying first order conditions is:

$$e_i^* = \frac{V_i^2 V_j}{(V_i + V_j)^2}.$$

How do optimal effort levels change as players' valuations change? The simplest effect of changing valuations is that Player i 's optimal effort level is monotonically increasing in her own valuation to winning the contest.³ As the contest prize becomes more valuable to Player i , she is willing to exert more effort to win the prize, regardless of Player j 's valuation. We call this the "own value" (OV) effect.

The effect of V_j on Player i 's optimal effort level, however, depends on the two players' *relative* valuations.⁴ When Player j values the prize more than Player i , increasing V_j decreases the marginal return to effort for Player i , which *decreases* i 's optimal effort. We call this the "getting deterred" (GD) effect. When Player j values the prize less than Player i , increasing V_j *increases* i 's optimal effort. As V_j increases, the marginal utility to effort, which helps Player i retain the prize she values so highly, also increases. We call this the "doing the deterring" (DD) effect.

Player i 's optimal effort level thus varies non-monotonically with Player j 's valuation. These two effects are akin to deterrence. The player with the higher valuation responds to increases in her opponent's valuation and subsequent effort levels with reciprocal increases in her own effort. The player with the lower valuation responds to increases in her opponent's valuation and subsequent effort levels with decreases in her own effort.

The three dimensions to the comparative statics of a value show that the strategic interaction between players is more complicated than a simple discouragement effect. Suppose the players start with equal valuations, and then i 's value increases and j 's decreases. This is one of the treatments in the contest experiment by Anderson and Stafford (2003). A comparison of effort levels under treatment (asymmetric) and control (symmetric values) conflates the three dimensions of the shock's effect on effort. The experimental protocol described below is designed to detect and decompose all three effects.

Until now, we have discussed only the effects of changing valuations but have not discussed

³Formally, $\frac{\partial e_i^*}{\partial V_i} = \frac{2V_i V_j^2}{(V_i + V_j)^3} > 0$.

⁴This is because the sign of $\frac{\partial e_i^*}{\partial V_j} = \frac{V_i^2 (V_i - V_j)}{(V_i + V_j)^3}$ depends on the $V_i - V_j$.

changing marginal costs to effort. However, increasing one player's valuation is isomorphic to decreasing their marginal costs of effort. The three effects identified above also obtain for changes to marginal costs of effort. The effect of decreasing c_i on the optimal effort of both players is the same as the effect of increasing V_i (Corchón, 2007).

Experimental Design and Procedures

We described the game to subjects as a “Lottery Contest Game” in which two players compete for a prize. We informed subjects that the prize would be worth different amounts to each player each time they played the game and that they would know their exact values and their opponent's values when making their decision. The decision was described in terms of “purchasing contest tickets,” with the probability of winning the prize equal to one's share of total tickets purchased in the round. Each ticket cost 1 point, and subjects received a fresh endowment of 1000 points in every round. They kept whatever portion of their endowment they didn't spend each round.

In addition to manipulating valuations, our design also manipulates the information available to the subjects: availability of feedback and the presence of a payoff calculator to test whether varying information or providing a computational aid might enhance subjects' strategic thinking and encourage choices closer to the equilibrium predictions. As described in greater detail below, we vary prize valuations across rounds of the game, vary information across parts of each session (within subjects), and vary the payoff calculator across sessions (between subjects).

Each experimental session was divided into two parts. At the beginning of Part 1, subjects received written instructions explaining the contest game (see the Appendix). After the experimenter read the instructions out loud, subjects took a brief comprehension quiz. In Part 1, subjects played 17 rounds of the two-player asymmetric contest game without feedback about who won the contest or about the other players' actions in any round. In every round, subjects were randomly matched with another player and were informed of each player's valuation of the prize.

In Part 2, subjects played another 17 rounds with feedback about the results of each contest they played. The information provided in the feedback screen included the effort level of each player, the probability that each player would win the prize given the chosen effort levels, which player won the contest, and each player’s payoff (denominated in points). Each subject’s individual history of play was also shown on the screen in each round of Part 2.

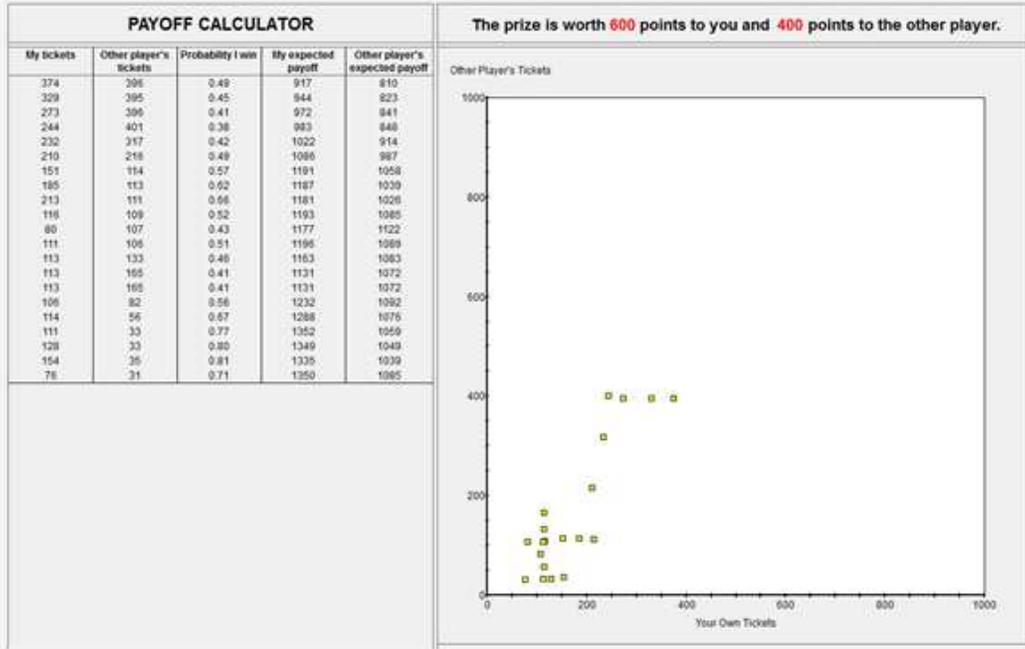
In the *Baseline* treatment sessions, subjects played each round of the asymmetric contest game without any computational aid. In the *Calculator* treatment, we provided subjects with a graphical interface in every round (shown in Figure 1). Crossing the feedback and calculator manipulations yields four conditions: *Baseline-No feedback* (BN), *Baseline-Feedback* (BF), *Calculator-No feedback* (CN), and *Calculator-Feedback* (CF).

The graphical interface for the calculator allows subjects to search the strategy space quickly and easily (see Figure 1). To do so, subjects clicked on a point in the white square on the right side of the screen (which represents the strategy space). Each time a subject clicked on the calculator, the pair of effort levels was displayed in a list on the left side of the screen along with the probability of winning and each player’s expected payoff. Each subject saw a list of all of their previous searches in that round, and at the beginning of every round the calculator was reset.

We selected eight distinct pairs of prize valuations (as shown in Table 1) to test the comparative static predictions. We refer to the set of valuations $S = \{200, 900\} \times \{300, 800\}$ as *single valuations*, and the set $D = \{400, 1800\} \times \{600, 1600\}$ as *double valuations* since each pair in the latter set is twice the value of a pair from the former. We also refer to prize values of 200, 300, 400, and 600 as *low values* and to prize values of 800, 900, 1600, and 1800 as *high values*.

If v_i is low, then increasing v_j from low to high generates the “getting deterred” effect (whereby e_i^* decreases), while if v_j is high, then increasing v_j from low to high generates the “doing the deterring” effect (whereby e_i^* increases). Each subject played each of these pairs twice, once as Player 1 and once as Player 2. We randomly generated a sequence of valuation pairs prior to the first session, with the order independent across parts 1 and 2, and held the sequence fixed across

Figure 1: Screenshot of graphical interface for payoff calculator



sessions. The last round in each part was a “zero value” round in which $v_i = v_j = 0$, which we included to measure whether any player has an unobserved preference for winning the contest (i.e., “joy of winning”).

We programmed the experiments in z-tree (Fischbacher, 2007) and conducted them in the [experimental laboratory at authors’ institution]. A total of 60 subjects participated in the experiment: 32 subjects participated in two sessions of the *Baseline* treatment and 28 subjects in two sessions of the *Calculator* treatment. At the end of each session, we randomly selected one round for payment and converted points to cash at a rate of \$1 per 75 points. Each session lasted less than an hour and a half, and subjects earned an average of \$21.40 (including a \$5 show-up fee).

The laboratory studies most closely related to ours investigate asymmetry across players and the so-called discouragement effect, where a stronger player (one with higher valuations, lower costs to effort, or better effort technology) induces the weaker player(s) to decrease effort. Anderson and Stafford (2003) vary participants’ costs to effort and find that costs are negatively associated with effort levels. Fonseca (2009), Anderson and Freeborn (2010), and Kimbrough,

Table 1: Valuations and Nash equilibrium predictions

	Valuations		Nash predictions		Expected payoffs	
	v_1	v_2	e_1^*	e_2^*	EU_1	EU_2
Single valuations	200	300	48	72	32	108
	200	800	32	128	8	512
	900	300	169	56	507	19
	900	800	224	199	253	177
Double valuations	400	600	96	144	64	216
	400	1600	64	256	16	1024
	1800	600	338	112	1014	38
	1800	1600	448	398	506	354

Sheremeta and Shields (2014) investigate games where players can have different effort technology. In these games, a unit of effort by a “strong” player has a greater marginal effect on her winning probability than a unit of effort from a weak player. They find that weak bidders generally exert less effort. Deck and Sheremeta (2012) analyze an experiment where a defender must defend against a sequence of attacks, and they vary the defender’s value to successfully defending all attacks. They find mixed support for a discouragement effect.

Experimental literature on contests also consistently identifies the phenomenon of “overbidding,” where players’ effort levels are much higher than the Nash equilibrium prediction. Dechenaux, Kovenock and Sheremeta (2014) observe that the degree of overbidding is sometimes high enough to give the players negative payoffs, meaning that they would have been better off not participating in the contest at all. There are likely many contributing factors to overbidding, such as if a player derives non-monetary utility from winning the contest or if the player simply makes mistakes.⁵

⁵Dechenaux, Kovenock and Sheremeta (2014) contains a section reviewing overbidding.

Results

Comparative Statics

We first analyze the results with respect to the comparative static predictions of the asymmetric contest game. To test the predictions, we estimate the following regression model of player i 's effort choice:

$$\text{Effort}_i = \alpha + \beta \text{High}_i + \gamma \text{GD}_i + \delta \text{DD}_i + \varepsilon_i$$

High is a dummy variable indicating that i 's value is high (where high and low values are defined as in the previous section), and the coefficient β measures the effect of increasing i 's own valuation. In theory, effort is increasing in one's own valuation, so we expect $\beta > 0$. GD is a dummy variable indicating that j 's value is high and i 's value is low, so the coefficient γ measures the getting deterred effect (how increasing j 's value affects i 's effort when i 's value is low). We expect from our theoretical analysis that $\gamma < 0$. Similarly, DD is a dummy variable indicating that j 's value is high and i 's value is high, so the coefficient δ measures the doing the deterring effect (how increasing j 's value affects i 's effort when i 's value is high). We expect $\delta > 0$.

Table 2 presents ordinary least squares regression estimates of this model for rounds with single valuations, and Table 3 presents the estimates for rounds with double valuations.⁶ The first column in each table presents estimates using data pooled across conditions (NN, NF, CN, CF), and the remaining columns present separate estimates for each condition. To account for within-subject dependence, we use robust standard errors clustered at the subject-level.

The first thing to note about our results is that subjects respond naturally to increases in their own valuations. Looking at the first column in Table 2, we see that subjects purchase about 104 tickets on average when both players' valuations are low and then purchase 205 more tickets when their own valuation increases. The exact magnitudes vary somewhat across conditions, but the

⁶This analysis excludes the zero valuation rounds.

Table 2: Evaluation of comparative static predictions (single valuations)

	Combined	NN	NF	CN	CF
	b/se	b/se	b/se	b/se	b/se
Own high	204.60** (11.74)	266.33** (21.46)	184.17** (17.37)	220.07** (26.08)	141.95** (19.03)
Getting deterred	-29.98* (11.67)	-55.91* (26.67)	-22.83 (17.57)	-10.36 (22.08)	-28.14* (11.15)
Doing deterring	-23.07+ (13.16)	-94.13** (29.79)	9.75 (21.73)	-43.63+ (22.78)	41.18+ (24.05)
Constant	104.44** (13.44)	129.14** (30.46)	87.30** (18.45)	104.61** (13.66)	95.64** (14.47)
Observations	960	256	256	224	224
R^2	0.30	0.33	0.29	0.32	0.32

OLS with standard errors clustered by subject

+ $p < .10$, * $p < .05$, ** $p < .01$

Table 3: Evaluation of comparative static predictions (double valuations)

	Combined	NN	NF	CN	CF
	b/se	b/se	b/se	b/se	b/se
Own high	292.81** (21.67)	368.14** (35.23)	292.45** (31.89)	280.48** (33.16)	219.45** (30.38)
Getting deterred	-69.91** (14.40)	-55.42+ (27.56)	-60.67* (23.81)	-87.45** (17.91)	-79.48** (15.89)
Doing deterring	-21.43 (21.17)	-148.14** (35.84)	76.55+ (40.33)	-51.39 (39.99)	41.36 (38.21)
Constant	189.73** (13.87)	206.38** (25.27)	155.69** (22.69)	227.73** (20.74)	171.61** (14.95)
Observations	960	256	256	224	224
R^2	0.37	0.33	0.46	0.35	0.45

OLS with standard errors clustered by subject

+ $p < .10$, * $p < .05$, ** $p < .01$

overall effect of increasing one's own valuation is consistently positive. As Table 3 shows, the same holds for double valuation rounds, with effort levels that are nearly doubled as equilibrium theory would predict: subjects purchase an average of 189 tickets when their valuations are low and an additional 293 tickets when their valuations are high.

The also results provide evidence for the getting deterred effect. In both Table 2 and Table 3, all of the estimated coefficients are negative, regardless of whether we pool the data or estimate the model separately for each experimental manipulation. The coefficients are also significant for the pooled single valuations and the pooled double valuations, and across the board for double valuations regardless of the availability of information or the calculator. The evidence for the getting deterred effect is somewhat weaker for single valuation rounds, as the coefficients in the NF and CN conditions are negative but not significant. Overall, consistent with the comparative static prediction, we find that when subjects' own valuations are low, they respond to increases in their opponent's valuation by reducing their effort.

The evidence for the doing the deterring effect is mixed, as subjects respond differently across information conditions. Pooling single valuation or double valuation rounds, we find that coefficient estimates are negative rather than positive (significant at the 0.10 level for single valuations and insignificant for double valuations). This appears to contradict the theoretical prediction. However, there are differences in the coefficients when we look at them separately by condition. When subjects do not receive feedback in Part 1 (the NN and CN conditions), they respond to increases in their opponent's valuation by decreasing their effort, and these effects are largest in the NN condition for both single and double valuations. In contrast, behavior is consistent with the theoretical predictions with feedback in Part 2, as subjects respond to their opponent's valuation with increases in effort. The coefficients for the NF and CF conditions are positive in both Table 2 and Table 3, and they are also marginally significant (at the 0.10 level) in the CF condition for single valuations and the NF condition for double valuations. The data suggest that the change in subjects' effort levels depend on their information: they decrease their effort without feedback while they increase

it with feedback.

Information Effects

How do the different treatments affect strategic behavior? As a simple assessment of the treatment effects, we first analyzed whether providing information or the payoff calculator promotes effort levels closer to the Nash equilibrium predictions. In this analysis, the dependent variable of interest is the percentage difference between effort and the equilibrium prediction:

$$\text{Pct. Difference} = \frac{\text{Effort} - \text{Prediction}}{\text{Prediction}}$$

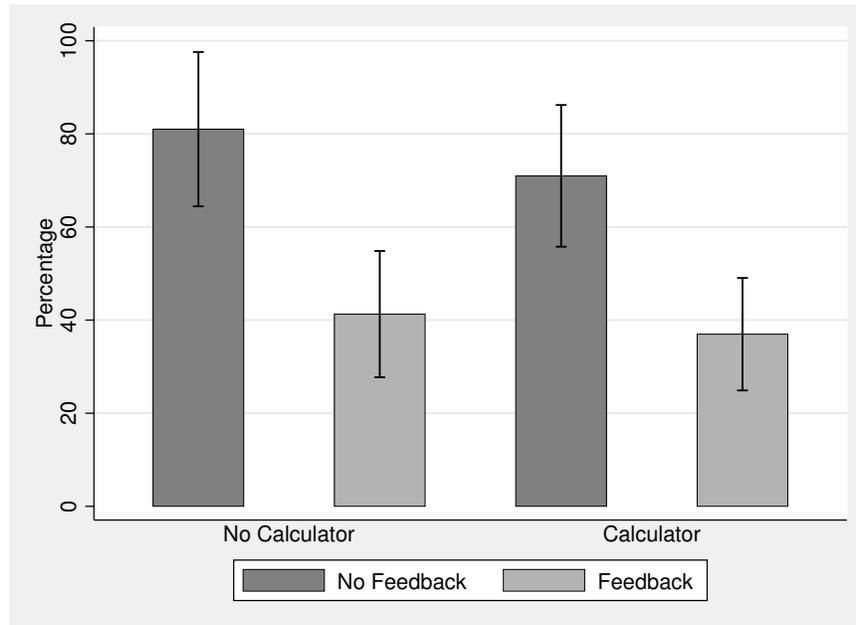
Lower values indicate behavior that is closer to the predictions while higher values indicate the opposite. We also pool the data across valuation pairs in each condition and exclude the zero valuation round from analysis.

Figure 2 presents the average difference by experimental manipulation. It shows a high level of over-effort (relative to the Nash prediction) in the baseline NN condition (81%). Introducing feedback has a substantial effect in the *Baseline* treatment, reducing over-effort by half (to 41.3%). Introducing the payoff calculator (holding the absence of feedback constant) appears to have only a slight effect on bringing effort closer to theoretical predictions (to 71%). However, the effect of introducing feedback in the *Calculator* treatment is substantial and similar in magnitude to the effect in the *Baseline* treatment, reducing the level of over-effort in half there as well (to 37%).

Table 4 presents OLS regression estimates of the treatment effects with robust standard errors clustered by subject, and the results reinforce our interpretation of Figure 2. Feedback has a significant effect on reducing over-effort, while the calculator has a small, statistically insignificant effect. These estimates do not change if we include predicted effort levels and a dummy variable for double valuations as controls.

Since the feedback treatment always occurs in later rounds, we show in the appendix that

Figure 2: Effort as percentage of Nash prediction



the feedback treatment effects are not simply an artifact of learning and increasing familiarity with the game over time. There is a discontinuity in respondents' behavior after the feedback treatment, compared with before. If learning explained the feedback effect, we would expect changes in features of behavior to be smooth around the period where the feedback treatment was administered. The appendix also contains results about the search behavior recorded from the subjects' use of the calculator. We explore various features of their searches, such as whether subjects search in areas that yield positive expected utility, the angle and direction of their searches over time, and the relationship between search quality and observed effort levels.

Strategic Sophistication

A large and growing body of work in behavioral game theory explains and organizes experimental subjects' departures from Nash equilibrium and variation in their behavior in terms of heterogeneous levels of strategic sophistication (Agranov et al., 2012; Arad and Rubinstein, 2012; Craw-

Table 4: Treatment effects and individual-level covariates

	(1)	(2)	(3)
	b/se	b/se	b/se
Feedback	-37.06** (8.40)	-37.06** (8.41)	-37.06** (8.42)
Calculator	-7.17 (21.89)	-7.17 (21.91)	-10.87 (19.25)
Nash prediction		-0.09+ (0.05)	-0.09+ (0.05)
Doubled valuations		-17.39* (7.53)	-17.39* (7.53)
Experience		-2.05** (0.76)	-2.05** (0.76)
Zero value effort			0.75** (0.07)
Male			-21.91 (26.02)
Risk avoidance scale			-34.99 (35.09)
Aggression scale			-57.73 (46.81)
Constant	79.68** (18.72)	120.25** (26.11)	142.29** (33.57)
Observations	1920	1920	1920
R^2	0.01	0.03	0.06

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$

+ $p < .10$, * $p < .05$, ** $p < .01$

ford, 2003; Gill and Prowse, 2012a). The basic idea is that some players are more strategic than others: more sophisticated subjects engage in higher orders of reasoning or form more accurate forecasts of others' behavior. The two most commonly used models, level-k and cognitive hierarchy (CH) models, were developed in the context of the "beauty contest" game, where players choose a number between 1-100 and are rewarded if they choose the value that is closest to $2/3$ times the average value the players chose (Nagel, 1995). The game is solvable by iterated reasoning, so players who engage in more steps of reasoning are described as having higher levels of strategic behavior. The level-k model describes the most naive, level-0, players as choosing randomly over the full support of possible actions. A level-1 player chooses the best response to a population of level-0 players. A level-2 player best responds to level-1 players, and so on. Players are classified by the level that is most consistent with their observed behavior, and players with higher levels are thought to have greater ability to anticipate their opponents' actions and incentives.

Experimental research has used a player's level as both an explanatory variable and an explanator. Most work in microeconomics emphasizes levels (Crawford, Costa-Gomes and Iriberri, 2013), so we will focus on the growing literature in political science that uses these methods. For example, Hafner-Burton et al. (2014) use a beauty contest game to measure participants' levels of strategic reasoning. They find that players with higher observed levels choose more strategically in an international negotiation and cooperation setting. Hafner-Burton, Hughes and Victor (2013) hypothesize that variation in the level of strategic thinking explains the difference between experimental results using college student convenience samples and those using more experienced policy elites. Bausch and Zeitzoff (2014) find variation in individuals' level of strategic thinking in a terrorist/counter-terrorist laboratory game. Loewen, Hinton and Sheffer (2015) find similar variation in a game about strategic voting in an election. Bassi (2015) analyzes how variation in electoral rules and candidate profiles can affect subjects' degree of strategic sophistication as measured by their level in a level-k model. Bassi and Williams (2014) analyze how variation in

financial incentives affects the level of subjects' strategic thinking. Minozzi and Woon (2013) find that a level- k model seems to fit non-equilibrium behavior in a communication game with competing experts.

We take a broad view of strategic sophistication, which we believe encompasses a variety of differences in strategic thinking beyond the number of steps of iterated reasoning. Thinking strategically means understanding incentives as well as understanding the behavior of others. One might be able to formulate the best response to an ideal opponent but misjudge how close one's opponent comes to that ideal. Conversely, one might accurately anticipate opponent's behavior but fail to recognize the optimal response.

Strategic reasoning also encompasses an individual's ability to react to changes in the strategic setting. In contrast to canonical beauty contest experiments in which the game is static, varying the prize values in our experiment means that subjects face a changing competitive landscape. Good strategic thinking also involves understanding comparative statics: how changes in the structure of the game affect incentives, behavior, and the feedback between these aspects of strategic interaction. Players with lower levels of strategic sophistication might understand how changes in their own prize values affect their own payoffs and incentives in a decision-theoretic sense (i.e., holding the behavior of others constant) but fail to ignore second order effects on the behavior of others. At higher levels of sophistication, players understand how others' prize values affect their opponents' behavior, thereby affecting their own behavior even if their own value of the prize remains constant.

Is there heterogeneity across respondents in the degree to which their behavior is consistent with the Nash comparative static predictions? Here, we focus only on the DD and GD effects, since there is strong support for the OV effects across subjects. We first construct a measure of whether individuals change their effort levels in the ways predicted by the DD and GD comparative statics. We leverage the fact that the subjects play each of their own valuation levels twice per treatment group. For example, subject i will play four rounds in which her value to the prize is 200, twice

in the first 16 rounds and twice in the last 16 rounds. For each of these rounds, her opponent’s valuation will be “high” once and “low” once. Looking again at Table 1, when player i has a value of 200, her opponent will sometimes have a low valuation of 300 and sometimes have a high value of 800. This allows us to assess how her own effort level changes as we vary her opponent’s effort level from low to high. When player i has a low valuation, moving her opponent from low to high should cause i to decrease her effort (the GD effect). When i has a high valuation, moving her opponent from low to high should cause i to increase her effort (the DD effect). We construct two binary variables, $\% \text{ Correct GD}$ and $\% \text{ Correct DD}$, in order to compare i ’s effort level as her opponent moves from low to high. These two variables equal 1 if the player’s observed behavior is in line with the comparative static prediction and 0 otherwise. Since each player plays 32 non-zero-valued rounds, this gives us 16 pairs of matched rounds per player. Next, we aggregate these data to the player-level and classify a player as having high *GD Consistency* if her behavior is consistent with the getting deterred prediction in more than half of these matched rounds and as having low consistency otherwise. We define high and low *DD Consistency* analogously.

Table 5 describes the distribution of subjects in terms of *GD Consistency* and *DD Consistency*. A large number of players, 26 out of 60, do not choose effort levels that are consistent with either prediction more than half the time. 5 players’ behavior is consistent with both predictions. In the margins, as expected above, more players’ behavior is consistent with the GD effect (28 out of 60) compared to the DD effect (11 out of 60).

Table 5: Subject-level consistency with opponent-value comparative statics

Doing deterring	Getting deterred		Total
	Low	High	
Low	43%	39%	82%
High	10%	8%	18%
Total	53%	47%	$N = 60$

Our analysis of comparative statics in Tables Table 2 and Table 3 provided some evidence that the Feedback treatment increases the level of strategic behavior. Recall that the own value

and getting deterred predictions received support that was generally consistent across treatments and valuations while the “doing the deterring” prediction received inconsistent support. In the Feedback treatments, the DD prediction receives the strongest support. Unlike in the calculator treatment rounds, the DD effect is signed correctly in the feedback rounds for both the high and low valuations. It is statistically significant in two of the four specifications.

The feedback treatment also has an effect on the individual level measures described above. Table 6 shows the classification of individuals based on their behavior during the feedback and non-feedback treatments. The feedback treatment increases the number of individuals whose behavior is consistent with each prediction, especially the DD effect. Without feedback, only 4 individuals’ behavior is consistent with the DD prediction, compared to 19 with feedback. The feedback treatment also increases the number of individuals whose behavior is consistent with the GD prediction as well, 22 without feedback compared to 26 with feedback.

Table 6: Consistency with opponent-value comparative statics by feedback treatment

<u>Without Feedback</u>			
Getting deterred			
Doing deterring	Low	High	Total
Low	60%	34%	94%
High	3%	3%	6%
Total	63%	37%	$N = 60$

<u>With Feedback</u>			
Getting deterred			
Doing deterring	Low	High	Total
Low	42%	26%	68%
High	15%	17%	32%
Total	57%	43%	$N = 60$

Conclusion

Many real-world situations are like contests where players value the prize differently. In wars, lobbying battles, campaigns, and other contests, the stakes of the contest can also change for one or more participants across issues, levels, or time. This paper has focused on how players respond to these shocks and to asymmetry. Under certain conditions, equilibrium theory predicts these shocks increase or decrease one or the other player's effort levels. Effort is important because the amount of effort exerted in a contest incurs opportunity cost—such effort and resources could have contributed to societal welfare, but did not. Understanding the effect of these shocks on effort is worthy of attention.

In a laboratory setting, we found support for some but not all of the comparative static predictions regarding asymmetry and valuation shocks. Intuitively, players increase their effort levels in response to positive valuation shocks, and vice versa. In terms of the predicted cross-player effects, we find that players “get deterred”: when player j values the prize less highly than player i , positive shocks to i 's valuation decreases j 's effort level. However, we find mixed support for “doing the deterring.” That is, theory also predicts that, when j values the prize more than i , increases in i 's valuation should increase j 's effort. We find that this prediction finds support only under treatments in which players have feedback about past rounds.

There is also significant heterogeneity across subjects in the degree to which their behavior is consistent with the Nash predictions, with some subjects displaying some, none, or all of the predictions. The feedback treatment, but not the calculator treatment, increases the strategic sophistication of the subjects' behavior. We speculate that feedback helps players better understand the strategic incentives when they respond to valuation shocks. This potential mechanism may stem from cognitive complexity because the own-value effects are easy to understand: players should exert more effort when the prize is worth more to them. The cross-player dynamics are subtler and more difficult to understand. It is possible that players have an easier time understand-

ing the “getting deterred” effect, but only understand the “doing the deterring” effect with greater experience, having observed past play via feedback.

We were surprised that the feedback treatment effects were generally stronger than the calculator effects, since the calculator is a much more powerful analytical tool than feedback, from the perspective of a player. With the calculator, the player can quickly learn a lot about the underlying payoff surfaces for herself and her opponent. If a player was willing to search the feedback space, iteratively finding best responses as in the strategic thinking used in level-k models, her search would eventually converge on the Nash equilibrium. Or, if she had a hypothesis about her opponent’s effort levels, she could use the calculator to find her best response. On the other hand, with feedback, the player only observes certain pieces of information regarding past play. It is possible that feedback is more experiential, causing the knowledge gained from feedback to be more meaningful or impactful than the information learned from the calculator.

This has implications for real world contests, where some contests are characterized by repeated interactions with the possibility for gaining experiential feedback. In international relations, for example, a large body of work focuses on enduring rivalries between countries (Diehl and Goertz, 2001). In American politics, some contests are over repeated topics, like appropriations, or regulatory battles over yearly quotas or rules. Our results suggest that, in these types of situations, where players have gained experiential feedback about the contest and their opponents over time, players’ behavior may be more consistent with Nash predictions. This may be more relevant than the degree to which the actors invest in analyzing a particular situation, with reports, technical consultation, spying, or other information-gathering mechanisms. Those efforts, which are qualitatively similar to the calculator treatment, may result in less strategic thinking and behavior than simple repetition and experience, which is qualitatively similar to the feedback treatment.

Our finding may also further explain divergence in behavior between undergraduate laboratory subjects and elites or politicians. The two groups may be similar in their ability to gather and process information about the payoff space, but the former group lacks the experiential learning

needed for more advanced strategic thought. Elites and politicians may behave more strategically or respond to asymmetry and shocks in more strategically sophisticated ways because they have experience with a particular contest.

Finally, our findings suggest the importance of broadening our conception of how strategic thinking varies across individuals. Level-k models based on iterative reasoning capture an important facet of strategic thinking, especially in games where iterative elimination of dominated strategies yields the Nash equilibrium. But individuals also vary in their ability to understand the underlying dynamics of interactions. They vary in their understanding of comparative statics: how changes to the features of the game pertaining to themselves and their opponents affect their incentives and choices. In many games that represent real world phenomena, this novel concept of strategic heterogeneity may more accurately describe an individual's degree of strategic sophistication.

References

- Abramowitz, Alan I. 1991. "Incumbency, campaign spending, and the decline of competition in US House elections." *The Journal of Politics* 53(01):34–56.
- Agranov, Marina, Elizabeth Potamites, Andrew Schotter and Chloe Tergiman. 2012. "Beliefs and endogenous cognitive levels: An experimental study." *Games and Economic Behavior* 75(2):449–463.
- Anderson, Lisa R and Beth A Freeborn. 2010. "Varying the intensity of competition in a multiple prize rent seeking experiment." *Public Choice* 143(1-2):237–254.
- Anderson, Lisa R. and Sarah L. Stafford. 2003. "An Experimental Analysis of Rent Seeking Under Varying Competitive Conditions." *Public Choice* 115(1-2):199–216.
- Arad, Ayala and Ariel Rubinstein. 2012. "The 11–20 Money Request Game: A Level- kReasoning Study." *The American Economic Review* 102(7):3561–3573.
- Bassi, Anna. 2015. "Voting systems and strategic manipulation: An experimental study." *Journal of Theoretical Politics* 27(1):58–85.
- Bassi, Anna and Kenneth C Williams. 2014. "Examining Monotonicity and Saliency Using Level- k Reasoning in a Voting Game." *Games* 5(1):26–52.
- Bausch, Andrew W and Thomas Zeitzoff. 2014. "Citizen Information, Electoral Incentives, and Provision of Counter-Terrorism: An Experimental Approach." *Political Behavior* pp. 1–26.
- Becker, Gary S. 1983. "A Theory of Competition Among Pressure Groups for Political Influence." *The Quarterly Journal of Economics* 98(3):371–400.
- Besley, Timothy and Torsten Persson. 2011. "The Logic of Political Violence." *The Quarterly Journal of Economics* 126(3):1411–1445.

- Bob, Clifford. 2012. *The Global Right Wing and the Clash of World Politics*. Cambridge University Press.
- Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics* 119(3):861–898.
- Corchón, Luis C. 2007. "The theory of contests: a survey." *Review of Economic Design* 11(2):69–100.
- Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review* 93(1):133–149.
- Crawford, Vincent P, Miguel A Costa-Gomes and Nagore Iriberri. 2013. "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications." *Journal of Economic Literature* 51(1):5–62.
- Davis, Christina and Krzysztof J Pelc. 2012. "Cooperation in Hard Times: Self-Restraint of Trade Protection." *Princeton University manuscript, June* .
- Dechenaux, Emmanuel, Dan Kovenock and Roman M. Sheremeta. 2014. "A survey of experimental research on contests, all-pay auctions and tournaments." *Experimental Economics* pp. 1–61.
- Deck, Cary and Roman M Sheremeta. 2012. "Fight or flight? Defending against sequential attacks in the game of Siege." *Journal of Conflict Resolution* 56(6):1069–1088.
- Diehl, Paul F and Gary Goertz. 2001. *War and peace in international rivalry*. University of Michigan Press.
- Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental economics* 10(2):171–178.

- Fonseca, Miguel A. 2009. "An experimental investigation of asymmetric contests." *International Journal of Industrial Organization* 27(5):582–591.
- Gill, David and Victoria L. Prowse. 2012a. "Cognitive ability and learning to play equilibrium: A level-k analysis."
- Gill, David and Victoria Prowse. 2012b. "A structural analysis of disappointment aversion in a real effort competition." *The American economic review* pp. 469–503.
- Goldstein, Judith and Lisa L. Martin. 2000. "Legalization, Trade Liberalization, and Domestic Politics: A Cautionary Note." *International Organization* 54:603–632.
- Grossman, Gene M. and Elhanan Helpman. 1994. "Protection for Sale." *The American Economic Review* 84(4):833–850.
- Hafner-Burton, Emilie M, Brad L LeVeck, David G Victor and James H Fowler. 2014. "Decision maker preferences for international legal cooperation." *International Organization* 68(04):845–876.
- Hafner-Burton, Emilie M, D Alex Hughes and David G Victor. 2013. "The cognitive revolution and the political psychology of elite decision making." *Perspectives on Politics* 11(02):368–386.
- Henn, Christian and Brad McDonald. 2014. "Crisis Protectionism: The Observed Trade Impact." *IMF Economic Review* 62(1):77–118.
- Iaryczower, Matias and Andrea Mattozzi. 2013. "On the nature of competition in alternative electoral systems." *The Journal of Politics* 75(03):743–756.
- Kimbrough, Erik O, Roman M Sheremeta and Timothy W Shields. 2014. "When parity promotes peace: Resolving conflict between asymmetric agents." *Journal of Economic Behavior & Organization* 99:96–108.

- Krueger, Anne O. 1974. "The Political Economy of the Rent-Seeking Society." *The American Economic Review* 64(3):pp. 291–303.
- Levitt, Steven D. 1994. "Using repeat challengers to estimate the effect of campaign spending on election outcomes in the US House." *Journal of Political Economy* pp. 777–798.
- Loewen, Peter John, Kelly Hinton and Lior Sheffer. 2015. "Beauty contests and strategic voting." *Electoral Studies* 38:38–45.
- Meirowitz, Adam. 2008. "Electoral contests, incumbency advantages, and campaign finance." *The Journal of Politics* 70(03):681–699.
- Minozzi, William and Jonathan Woon. 2013. "Lying aversion, lobbying, and context in a strategic communication experiment." *Journal of Theoretical Politics* 25(3):309–337.
- Myerson, Roger B. 2008. "The autocrat's credibility problem and foundations of the constitutional state." *American Political Science Review* 102(01):125–139.
- Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review* 85(5):1313–1326.
- Nordås, Ragnhild and Nils Petter Gleditsch. 2007. "Climate change and conflict." *Political Geography* 26(6):627–638.
- Ross, Michael L. 2004a. "How Do Natural Resources Influence Civil War? Evidence from Thirteen Cases." *International Organization* 58(01):35–67.
- Ross, Michael L. 2004b. "What do we know about natural resources and civil war?" *Journal of peace research* 41(3):337–356.
- Salehyan, Idean. 2008. "From climate change to conflict? No consensus yet." *Journal of Peace Research* 45(3):315–326.

- Sell, Susan K. and Aseem Prakash. 2004. "Using Ideas Strategically: The Contest Between Business and NGO Networks in Intellectual Property Rights." *International Studies Quarterly* 48(1):143–175.
- Serra, Gilles. 2010. "Polarization of what? A model of elections with endogenous valence." *The Journal of Politics* 72(02):426–437.
- Slantchev, Branislav L. 2010. "Feigning Weakness." *International Organization* 64(03):357–388.
- Stahl, D.O. and P.W. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10(1):218–254.
- Svolik, Milan W. 2009. "Power sharing and leadership dynamics in authoritarian regimes." *American Journal of Political Science* 53(2):477–494.
- Tullock, Gordon. 1967. "The Welfare Cost of Tariffs, Monopolies and Theft." *Western Economic Journal* 5(3):224–232.
- Welch, Susan and John R Hibbing. 1997. "The effects of charges of corruption on voting behavior in congressional elections, 1982–1990." *The Journal of Politics* 59(01):226–239.

Appendix for: How Hard to Fight? Strategic Sophistication in an Asymmetric Contest Experiment

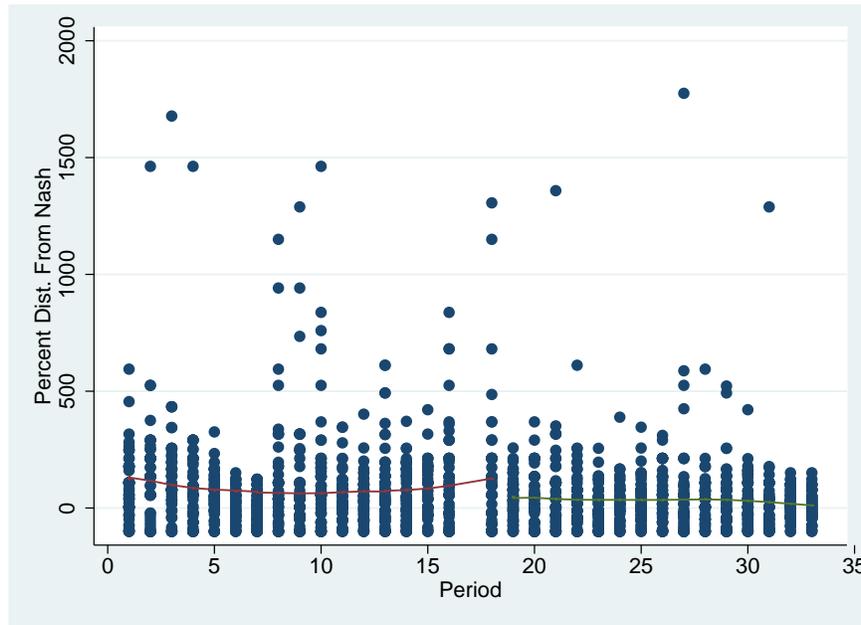
October 29, 2015

Feedback Effect or Learning Over Time?

One possible concern is that the effect of the feedback treatment can be attributed to learning as opposed to feedback. The feedback treatment is administered after 17 rounds of play, so observing that participants' efforts are closer to Nash in rounds 18-34, under the feedback treatment, may be an artifact of the experience or learning that took place in early rounds, as opposed to the treatment. To assess this, we look for a discontinuity in behavior before and after the feedback treatment. If learning explains the change in outcomes, then we should not see a discontinuity. The rate at which behavior converges towards Nash predictions should be steady before and after the treatment. If there is a jump, and behavior gets most closer to Nash predictions after the treatment, then this would suggest that the treatment effect is not an artifact of learning.

Figure A1 shows the percent distance from Nash predictions by period, with Lowess smoothers before and after the feedback treatment. Note that the treatment begins in Round 18, but since the feedback is only provided after participants choose their effort levels, the treatment is administered *after* they make their Round 18 choice. That is why the left side Lowess line includes the efforts

Figure A1: Percent Distance from Nash by Round

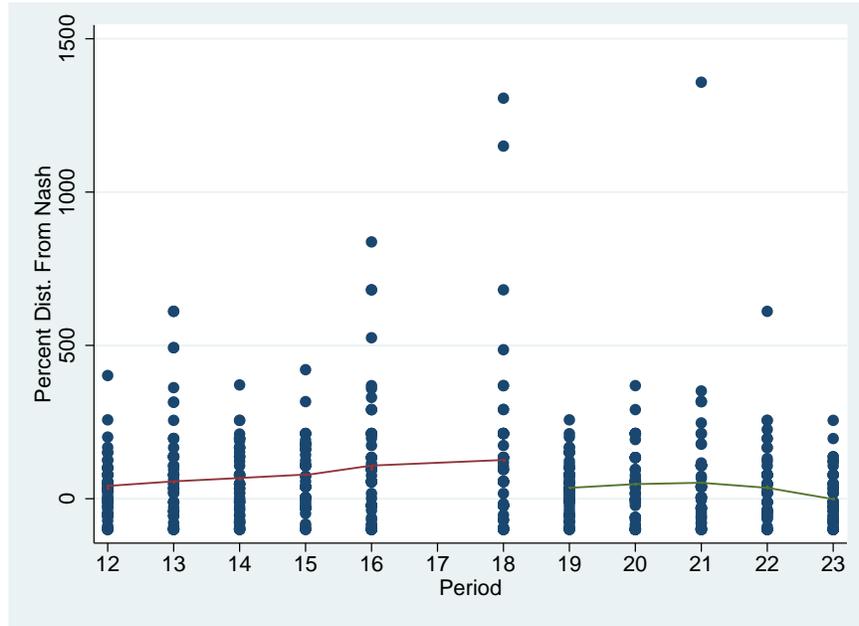


from Round 18. There is a slight decline in distance from Nash predictions over time, but there is a distinct jump downwards after the feedback treatment is administered. This jump is also apparent in Figure A2 which zooms in on the break point, only including Rounds 12-23. The distinct break supports the argument that the feedback treatment effect is not simply an artifact of learning over time.

Variation in Search Quality

Each time subjects used the payoff calculator, we stored the results. Although we found that the availability of the payoff calculator has no effect on overall effort levels, we can use the calculator data to analyze patterns of search behavior to investigate heterogeneity in strategic sophistication. For example, subjects might guess how many tickets their opponent will buy and then use the calculator to compute their own best response. Strategically sophisticated subjects might then try to calculate their opponent's best response to their effort level, iterating this process in a way

Figure A2: Percent Distance from Nash by Round, Plus/Minus Five Round Window



that traces out Cournot-style best response dynamics. With sufficient iterations, such subjects would be able to identify the Nash equilibrium of the asymmetric contest game. Strategically naive subjects might instead use the calculator to check the probability of winning or the expected payoffs associated with a given number of contest tickets, perhaps searching for an effort level until it exceeds an arbitrary threshold in a manner that reflects satisficing.

Motivated by these ideas, we construct several measures of search quality. For a set of minimal measures of search quality, we code whether each click or guess yields net positive expected utility, EU_i , relative to purchasing 0 tickets and ensuring a payoff of 1000 points. Own Positive indicates whether $EU_i > 1000$ for subject i (the subject using the calculator), Opponent Positive indicates whether the guess yields positive expected utility $EU_j > 10000$ for i 's opponent j , and Both Positive indicates searches where both $EU_i > 1000$ and $EU_j > 1000$.

Another measure of search quality relates to the direction of search. Let $g_k = (e_{ik}, e_{jk})$ denote subject i 's k -th guess in any given round. The direction of search refers to the angle of the difference vector $\Delta g = g_{k+1} - g_k$, which we measure in degrees (from 0° to 360°). If a subject searches

the strategy space by holding the opponent's effort constant $e_{j,k} = e_{j,k+1}$ while varying her own effort $e_{i,k} \neq e_{i,k+1}$, the direction of search will be *horizontal*. Conversely, if a subject holds her own effort constant $e_{i,k} = e_{i,k+1}$ while varying her guesses about her opponent's effort $e_{j,k} \neq e_{j,k+1}$, the direction of search will be *vertical*. Horizontal searches reflect a subject's attention to her own payoffs, which is individually rational in the sense of maximizing one's own payoffs, while vertical searches reflect attention to her opponent's payoffs and reflect strategic rationality in the sense of forming rational expectations about opponent behavior. We allow for two levels of error tolerance in how we classify horizontal and vertical searches, with a relatively narrow tolerance of $\pm 10^\circ$ and a wider tolerance of $\pm 22^\circ$. We then code each guess after the first ($k > 1$) as horizontal, vertical, or diagonal (neither horizontal nor vertical).

We find that the quality of subjects' searches according to these measures tends to be fairly poor. Table A1 describes the averages for our measures of search quality along with the distance between each guess and the total number of guesses. We present the overall means, the subject-period level means, and the subject-level means. The results do not differ much by level of aggregation.

According to our positive expected payoff measures, at most half of subjects' searches in the *Calculator* treatment can be classified as minimally rational. 53.1% of guesses involve positive expected values for the subject's own payoffs and 50.8% of guesses involve positive expected values for their opponent's payoffs. However, fewer than one-third of guesses (32%) involve positive expected payoffs for both the subject and their opponent. While we would expect to see that initial searches within a period yield net negative expected payoffs, we also thought that minimally rational search behavior would move quickly towards areas of the strategy space where both players receive positive expected utility. The prevalence of negative expected payoff guesses suggests to us that most searches are of low quality.

We also find that horizontal and vertical searches comprise half of the guesses entered into the calculator. While we might expect some searches to be diagonal, systematic guesses along one of

Table A1: Measures of search quality

	Total		Subject-period		Subject	
	Mean	N	Mean	N	Mean	N
Own Positive	.531	5067	.595	511	.524	28
Opponent Positive	.508	5067	.534	511	.506	28
Both Positive	.320	5067	.350	511	.316	28
Horizontal ($\pm 10^\circ$)	.311	4556	.334	452	.308	28
Horizontal ($\pm 22^\circ$)	.411	4556	.425	452	.405	28
Vertical ($\pm 10^\circ$)	.202	4556	.235	452	.195	28
Vertical ($\pm 22^\circ$)	.287	4556	.306	452	.266	28
Distance	142.9	4556	155.5	452	157.6	28
Searches	–	–	9.9	511	181.0	28

the dimensions to search for a player's best reply appear to be rare. Searches along one dimension also tend to be horizontal (31% of all searches using the 10° tolerance) rather than vertical (20%), which suggests that subjects tend to focus on their own payoffs rather than their opponents. This may reflect a failure of subjects to engage in any kind of meaningful strategic reasoning.

To assess whether the quality of search affects behavior in the asymmetric contest game, we estimate several regression models with our search measures as right-hand side variables. The independent variables are calculated as subject-period level averages.¹

We use three dependent variables. The first is simply the subject's observed effort level for that round. The second is the difference between i 's observed effort e_i and their empirical best response $b_i(e_j)$, where the empirical best response is a function of the opponent's actual effort e_j :

$$b_i(e_j) = \max \{ \sqrt{v_i e_j} - e_j, 0 \}$$

This quantity describes the optimality of the subject's play with respect to what her opponent actually chose. Negative coefficients indicate that that feature of search behavior was associated with the subject choosing an effort level that was more optimal, given the actual effort level of her

¹The search measures are rescaled so that they are percentages. Distance is also rescaled so that 10 tickets equals 1 unit. Time is measured in seconds from the beginning of the round until a choice is made.

opponent. The third dependent variable is the distance between the subject's effort level and the Nash equilibrium prediction for her effort, given the pair of valuations. The models are estimated using OLS with subject-level robust standard errors. We include the Nash equilibrium prediction as a control variable to account for differences in valuations across periods. For each dependent variable, the first column corresponds to rounds with no feedback (marked "NF") and the second column corresponds to rounds with feedback (marked "F").

Table A2: Effect of Searching on Effort Features

	Effort Level		Dist. from Emp. Best Resp.		Dist. from Nash Effort	
	(1)	(2)	(3)	(4)	(5)	(6)
	NF	F	NF	F	NF	F
Nash prediction	1.13** (0.13)	1.04** (0.09)				
Own positive	0.84* (0.35)	1.04** (0.34)	1.03** (0.37)	0.94* (0.44)	-0.10 (0.41)	0.11 (0.32)
Other positive	-0.36 (0.28)	-0.27 (0.23)	-0.24 (0.29)	-0.49 (0.30)	-0.10 (0.30)	-0.04 (0.42)
Both positive	-1.74** (0.30)	-1.57** (0.47)	-1.46** (0.37)	-1.30+ (0.67)	-0.93** (0.31)	-1.13+ (0.56)
Distance	1.32 (0.89)	1.15 (1.06)	1.71 (1.01)	1.13 (1.25)	1.66 (1.49)	2.43 (1.73)
Horizontal ($\pm 10^\circ$)	0.21 (0.44)	0.34 (0.51)	0.04 (0.50)	0.40 (0.59)	0.58 (0.58)	1.16* (0.43)
Vertical ($\pm 10^\circ$)	0.63 (0.39)	0.05 (0.49)	0.25 (0.43)	0.56 (0.76)	0.45+ (0.26)	0.40 (0.36)
Total number	-2.24 (1.45)	0.95 (1.67)	-2.06 (1.51)	0.84 (1.70)	-2.01 (1.39)	-0.09 (1.51)
Total time	1.43+ (0.77)	2.12* (0.95)	1.96* (0.90)	1.69+ (0.99)	0.97 (0.66)	1.24+ (0.66)
Constant	22.70 (42.80)	-33.37 (29.49)	50.63 (49.45)	18.94 (28.87)	41.20 (33.64)	-9.70 (18.48)
Observations	448	448	448	448	448	448
R^2	0.45	0.48	0.13	0.06	0.11	0.07

+ $p < .10$, * $p < .05$, ** $p < .01$

Table A2 shows the results. Overall, the effects of searching on effort features are consistent across the different effort features and across the Feedback and No Feedback conditions. The

strongest result is that a greater number of searches in the “both positive” region yielded lower effort levels and effort levels that were closer to both the Nash prediction and empirical best response. A greater number of searches in the “Opponent positive” region also yielded lower efforts and efforts closer to Nash predictions and empirical best responses, though these results are not significant. Players who engage in more rational search of their opponent’s strategies (as indicated by Other positive or Both positive) tend to choose lower levels of effort. There are not strong effects for the distance searched, number of searches, or the directionality of the searches. This suggests that where the subjects searched was more important than how many clicks or in what direction they searched.

Interestingly, a greater number of searches in the “Own positive” region tended to increase effort levels and also increase the distance between effort and the Nash prediction. This could be because respondents ascribed too low of an effort level to their opponent. Effort is increasing in total searchtime. The more time a player spent searching the space, the higher the eventual effort level that is chosen and the greater the distance to the Nash prediction and empirical best response.

Instructions

General Information

This is an experiment on the economics of strategic decision-making. XXX has provided funds for this research.

You will be paid in cash for your participation, and the exact amount you receive will be determined during the experiment and will depend partly on your decisions, partly on the decisions of others, and partly on chance. You will be paid your earnings privately, meaning that no other participant will find out how much you earn. These earnings will be paid to you at the end of the experiment along with the \$5 participation payment.

Pay attention and follow the instructions closely, as we will explain how you will earn money and how your earnings will depend on the choices that you make. Each participant has a printed copy of these instructions, and you may refer to them at any time.

If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you. Please do not talk, exclaim, or try to communicate with other participants during the experiment. Also, please ensure that any phones or electronic devices are turned off and put away. Participants intentionally violating the rules will be asked to leave and may not be paid.

Parts, Rounds, and Matching

This experiment consists of several parts. We will explain the instructions for each part before beginning that part. In each part, you will make decisions in one or more rounds.

In every round you will be **randomly matched** with one other participant. You will not know the identity of the other participant you are matched with in any round, and your earnings for each round depend only on your action in that round and the action of the participant you are matched with in that round.

Your earnings during each round are denominated in points, which we will convert to cash at a rate of \$1 for every 75 points. We will **randomly select one round to count** for payment from the entire session, and each round is equally likely to be selected. The points you receive in that round will be used to calculate your payment for the experiment. You should think of each round as a separate decision task.

Appendix: Instructions for Calculator Treatment

Part 1. Lottery Contest Game

In each round, you and the other participant you are matched with will compete for a prize. This prize will be worth X points to you and Y points to the other player. These amounts may be different in every round, and during the round both you and the other player will know exactly what the prize is worth to each of you.

You will compete for the prize by purchasing “contest tickets.” At the beginning of each round, you have 1,000 points. You can use these points to purchase contest tickets at a cost of 1 point per ticket. You can purchase up to 1,000 of these tickets. Any points you do not spend on contest tickets will be added to your point balance for the round.

Your payoff will be the number of tickets you keep plus, if you win, the value of the prize. If you buy T tickets and the prize is worth X points to you, then:

$$\text{Your payoff if you win} = X + 1000 - T$$

$$\text{Your payoff if you do not win} = 1000 - T$$

For example, suppose you buy 300 contest tickets and the prize is worth 600 points to you. Thus, you kept 700 points from your original 1,000 points. If you win the prize, then you would earn 600 points from the prize plus the 700 points you kept for a total earning of 1,300 points for the round. If you do not win the prize, then you would earn 700 points for the round. Of course, this is just one example of how to compute your possible earnings.

The winner of the prize is determined by a lottery contest. The lottery contest works as follows. As soon as everybody has chosen how many contest tickets to buy, the computer will randomly select one winning ticket (separately for each group) to determine whether you or the other player wins the prize. Your chance of winning the prize in the round depends on how many contest tickets you buy and how many contest tickets the other player buys. More specifically, your chance of winning is equal to your share of the total tickets bought in that round:

$$\text{Chance of winning prize} = \frac{\text{Your tickets}}{\text{Your tickets} + \text{Other player's tickets}}$$

For instance, if you and the other player each bought the same number of contest tickets, each of you has a 50 percent share of the lottery tickets and therefore a 50 percent chance of winning. If you buy twice as many contest tickets as the other player, you have two-thirds of the contests tickets (and therefore a two-thirds chance of winning) while the other player has a one-third share of tickets (and a one-third chance of winning).

Thus, your chances of winning the prize increase with the number of contest tickets you buy. Conversely, the more contest tickets the other player buys, the higher the probability that the other player wins. If only one player buys contest tickets, then that player will win the prize for sure. If nobody buys any contest tickets, no contest takes place and no one wins the prize.

Appendix: Instructions for Calculator Treatment

After everyone chooses how many tickets to buy in each round, we will proceed to the next round. You will not find out the results from any round of Part 1 until all rounds of Part 1 are completed.

Payoff Calculator

In every round, you will have access to a payoff calculator to help you make your decision (as shown in the first picture on the next page). To use the payoff calculator, click on a point inside the white square on the right side of the screen. You can think of the coordinates of the point you click as guesses about the possible amounts of tickets that you and your opponent might buy. The x-coordinate (along the horizontal dimension) corresponds to the number of tickets you might buy for yourself. The y-coordinate (along the vertical dimension) corresponds to the number of tickets you think the other player might buy.

For each time you click inside the white square, the results of the calculation will appear in a list on the left side of the screen as follows. The first two columns show you the numbers of tickets you entered into the calculator. The rest of the columns (from left to right) show you three useful quantities calculated for you:

- Your probability of winning the prize
- The “expected value” of your payoff
- The “expected value” of the other player’s payoff

The expected values describe the average number of points you might receive based on the tickets purchased in that round. The expected values are calculated using the following formula:

$$\text{Expected Value} = (\text{Prob. of win})(\text{Points from win}) + (\text{Prob. of loss})(\text{Points from loss})$$

The calculator will show you the results of all the calculations you made in that round, and you should use it as often as you need to before making a decision.

When you are ready to purchase contest tickets, click on the “Submit Decision” button in the bottom-right of the screen. When you click this button, you will see the Decision Input area on the right side of your screen (as shown in the second picture). This button will appear 20 seconds after the round begins so that you have some time to use the calculator. Note that you can also return to the calculator input box from the Decision Input screen and continue to use the payoff calculator as often as you like until you submit your decision. There is no time limit for using the calculator. To purchase your tickets in the Decision Input screen, enter a number in the box on the right side of the screen and then click on the red “Buy Tickets” button.

Appendix: Instructions for Calculator Treatment

Sample screens

PAYOFF CALCULATOR				
My tickets	Other player's tickets	Probability I win	My expected payoff	Other player's expected payoff

The prize is worth points to you and points to the other player.

Other Player's Tickets

Your Own Tickets

Click "SHOW DECISION" when you are ready to submit your decision.

PAYOFF CALCULATOR				
My tickets	Other player's tickets	Probability I win	My expected payoff	Other player's expected payoff

YOUR CONTEST DECISION

You are participant 2

The prize is worth points to you and points to the other player.

You have **1000 points** that you can spend to buy contest tickets.
Each ticket costs **1 point**

HOW MANY TICKETS WOULD YOU LIKE TO BUY?
(You can buy any whole number of tickets from 0 to 1000.)

Click "SHOW INPUT" to see the calculator input box.

Appendix: Instructions for Calculator Treatment

Instruction Quiz

Before we begin the experiment we would like you to answer a few questions to make sure you understand how the lottery contest game works. Please answer these questions on your computers. You will receive immediate feedback once you answer all of the questions. We will then begin the experiment when everyone has answered these questions.

1. Suppose the prize is worth 700 to you. If you purchase 100 tickets, how many points will you earn if you win the prize?
 - a. 600
 - b. 900
 - c. 1600
 - d. 1700

2. If the prize is worth 400 to you and you purchase 200 tickets, how many points will you earn if you do not win the prize?
 - a. 200
 - b. 400
 - c. 600
 - d. 800

3. If you purchase 100 tickets and the other player purchases 400 tickets, what is your chance of winning the prize?
 - a. $100 / 400$
 - b. $300 / 400$
 - c. $100 / 500$
 - d. $400 / 500$

4. If you purchase 300 tickets and the other player purchases 100 tickets, what is your chance of winning the prize?
 - a. $100 / 300$
 - b. $200 / 300$
 - c. $100 / 400$
 - d. $300 / 400$

Appendix: Instructions for Calculator Treatment

Part 2. Lottery Contest Game with Feedback

You will play the Lottery Contest Game in Part 2 exactly the same way you did in Part 1. The only difference is that between rounds, you will find out which player won the contest, how many tickets the other player purchased, and the number of points you earned during the round. During the round, you will also be able to view the results of all previous rounds you played, and you can switch between this history and the payoff calculator when making your decision in each round.