# The Compromise Game: Two-sided Adverse Selection in the Laboratory *

**Juan D. Carrillo**
*University of Southern California and CEPR*

**Thomas R. Palfrey**
*California Institute of Technology*

First version: May 2006

This version: November 23, 2006

*Part of this research was conducted while the first author was visiting Caltech. The hospitality of the hosting institution is greatly appreciated. We also gratefully acknowledge the financial support of the National Science Foundation (SES-0079301, SES-0450712, SES-0094800), The Princeton Laboratory for Experimental Social Science, and The Princeton Center for Economic Policy Studies. We thank seminar audiences at Collège de France, Princeton University, Universitat Autonoma de Barcelona, University of Pennsylvania and ESA meeting (Tucson) for comments, and Shivani Nayyar, Stephanie Wang, and Rumen Zarev for research assistance. Address for correspondence: Juan D. Carrillo, Department of Economics, University of Southern California, Los Angeles, CA 90089-0253, `<juandc@usc.edu>` or Thomas R. Palfrey, The Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, `<trp@hss.caltech.edu>`.

## Abstract

*We analyze a game of two-sided private information characterized by extreme adverse selection, and study a special case in the laboratory. Each player has a privately known "strength" and can decide to fight or retreat. If either chooses to fight, there is a conflict; the stronger player receives a high payoff and the weaker player receives a low payoff. If both choose to retreat, conflict is avoided and they each receive an intermediate payoff. The only equilibrium in both the sequential and simultaneous versions of the game is for players to always fight. In our experiment, we observe among other things (i) frequent retreat, (ii) no evidence of learning, and (iii) different behaviors between first, second and simultaneous movers. We explore several models in an attempt to explain these anomalous choices, including quantal response equilibrium, cognitive hierarchy, and cursed equilibrium.*

JEL classification: O24, O26.

Keywords: two-sided private information, adverse selection, laboratory experiment, behavioral game theory, quantal response equilibrium, cognitive hierarchy, cursed equilibrium.

# 1 Introduction

One of the major insights from theoretical research in information economics is that profitable agreements may be severely impeded by private information, and can even dry up completely. This was nicely illustrated in Akerlof's (1970) famous market for lemons example and studied in further detail by Myerson and Satterthwaite (1983) in a context of optimal contracting with two-sided private information. More generally, no-trade theorems (Milgrom and Stokey (1982), Morris (1994)) show that rational, expected utility maximizing, Bayesian economic agents will not trade with each other on the basis of private information alone.

In this paper, we study an environment where exchange, or other mutual agreements that have exchange-like features, is impeded for reasons of private information. A all-too-familiar example, war, illustrates the problem. Suppose there are two nations, either of which would be better off if conquering the other nation, compared to peaceful coexistence, and would be worse off being conquered. If there is a war, whichever country is strongest conquers the other one. The leader of each nation knows its own military strength but knows only the probability distribution of the other nation's strength. Each nation chooses to either "attack" or "not attack". They remain in peaceful coexistence if both choose not to attack, and a war ensues otherwise. If one formalizes this problem, the equilibrium always results in war. This would be true, for example, even if the benefits of winning the war were only slightly greater than the peace benefits, and the cost of losing a war were enormous. The logic is much like the unravelling argument in adverse selection games. In deciding whether to attack or not, optimal decision making requires the agents to condition on their opponent choosing "not attack". Because weaker opponents are the ones who do not attack, this conditioning will lead stronger opponents to attack. Therefore, there will be a marginal strength level which is indifferent between peace and forcing a war. But this calculus will lead the opponent's marginal non-attackers to attack, and so forth. The only equilibrium is for the marginal strength type to be the weakest type. As developed in section 2.1, the same logic applies to other situations where parties with conflicting goals and private information can negotiate an agreement: litigations, electoral debates, firm competition, and other situations.

We report here an experiment that explores behavior in several variations of this two-sided asymmetric information environment in the laboratory. In all the variations, the equilibrium outcome is the same: a war ensues with probability 1. We obtain several strong results. First, as predicted by the theory, agents roughly follow an individual cutpoint rule, where attack is selected if and only if strength is above a certain threshold. However, instead of this threshold being at (or at least close to) the minimum strength, we find that players generally use cutpoints in an intermediate range. As a result, peaceful outcomes, or compromises, occur with surprisingly high frequency – nearly 25% of the time in some sessions. Second, attack rates are affected by the treatment variables. In the sequential version and conditional on strength, second movers are significantly more likely to attack than first movers. Thus, second movers anticipate that not attacking is an indicator that the first mover had a relatively low strength. Also, in both the sequential and simultaneous treatments, agents are less likely to attack the higher the payoff under peace. Third, these results are robust with respect to experience, suggesting that there is little evidence of learning. It is puzzling that a subject who has played repeatedly as a second mover does not carry over his knowledge when becoming a first mover.

We then apply several recent theories of imperfect rationality in games to analyze the data and see if the insights from these alternative theories can account for our findings. The three approaches we explore are *equilibrium stochastic choice*, *levels of strategic sophistication*, and *naïve beliefs*. The specification of our models for these three approaches are, respectively, the logit specification of Quantal Response Equilibrium theory (QRE); the Poisson specification of the Cognitive Hierarchy model (CH); and a stochastic choice version of Cursed Equilibrium (CE). We find that all three models capture the main features of the data in remarkably similar ways. The estimated parameters for each model are relatively constant across treatments. However, there are some subtle differences in the predictions, that lead to differences in the fit of the models. The QRE model captures the tendency of the second mover to attack more often than the first mover. The CH and CE models capture the aggregate tendency of players to attack with probability close to 1 when their strength is sufficiently high and with probability close to 0 when their strength is sufficiently low. Not surprisingly then, the best fit is obtained with a combination of models.

# 2 The theoretical model

We analyze the incentives of agents to compromise when they have conflicting objectives and asymmetric information. To this end, we study a class of games that have unique Nash equilibrium outcomes in which a compromise is never reached. We start with some concrete examples that illustrate the conflicts we have in mind.

## 2.1 Some introductory examples

Consider two agents who must decide whether to split a surplus in a prespecified manner (compromise) or try to reap all the benefits for themselves (no-compromise). Both agents have private but imperfect information about their likelihood of obtaining the benefits if they do not compromise and, possibly, its value also. The ex-post sum of utilities may be higher or lower under compromise than under no-compromise.

A myriad of examples fit this general description. In a litigation, the defendant may offer a settlement to the plaintiff which can be accepted or not. Both parties have private knowledge of the strength of their case and the bias of the jury. In a military dispute, there is an existing division of the territory that can be challenged by either side depending on their privately known military strength. When a conflict breaks out, the winner annexes the territory of the loser. In an electoral campaign, each candidate can drive its rival into a public debate. If a debate takes place, the true talent of the contenders is revealed to voters and affects their probability of winning the election. If there is no debate, voters must rely on expected talents. In a product market competition, firms offering horizontally differentiated products may start an R&D race. The winner captures the entire market and the probability of winning is proportional to the privately known quality of their research department. Alternatively, firms can avoid the race and split the market. In all these cases, there are only two possible outcomes: settlement, peace, no debate, no race vs. trial, war, debate, race. The first outcome needs the agreement of both agents whereas each agent can unilaterally force the second outcome. Payoffs depend on the state of the world, which is not realized (or revealed) until after all players have acted. The total surplus of agents may vary across outcomes: litigations and wars are typically costly whereas electoral debates are neutral for the parties involved.[1] Last, the utility of agents under the different outcomes may depend on the private information parameters, but they are exogenously given. Payoffs under agreement are typically determined by the status quo situation whereas payoffs under no-agreement are typically determined by a winner-takes-all rule.

## 2.2 A formalization of the problem

We formalize the problem as follows. Denote by $s_i \in S_i$ and $s_j \in S_j$ the privately known "strength" of agents $i$ and $j$, with $i, j \in \{1, 2\}$ and $i \neq j$ (case strength, military capacity, politician's talent, research quality). These values are drawn from continuous and commonly known distributions $F_i(s_i \,|\, s_j)$ possibly different and possibly correlated. For technical convenience, we assume strictly positive densities $f_i(s_i \,|\, s_j)$ for all $s_i$ and $s_j$. Agent $i$ chooses action $a_i \in A = \{\rho, \phi\}$, where $\rho$ stands for "retreat" and $\phi$ for "fight". If $a_1 = a_2 = \rho$, there is compromise (settlement, peace, no debate, no race) and the payoff of agent $i$ is $\beta_i(s_1, s_2)$. Otherwise, there is no-compromise (trial, war, debate, race) and the payoff of agent $i$ is $\alpha_i(s_1, s_2)$ if $s_i > s_j$ and $\gamma_i(s_1, s_2)$ if $s_i < s_j$, with $\alpha_i(s_1, s_2) > \beta_i(s_1, s_2) > \gamma_i(s_1, s_2)$ for all $i, s_1, s_2$. Note that, ex-post, a compromise is always beneficial for one agent and detrimental for the other. The pair of strengths $(s_i, s_j)$ determines the winner and the loser. Payoffs under compromise and no-compromise are exogenously given, although they may be unknown at the time of making the decision if they depend on $(s_1, s_2)$. Last, the socially efficient action may be compromise or no-compromise or it may even be a zero-sum game: $\alpha_i(s_1, s_2) + \gamma_j(s_1, s_2) \gtreqless \beta_i(s_1, s_2) + \beta_j(s_1, s_2)$ for all $s_i > s_j$.

## 2.3 The optimal strategy

Given this structure, we can analyze the Perfect Bayesian Equilibrium (PBE) for the sequential version of the game. We have the following result.

**Proposition 1** *In all PBE of the game, the outcome is "no-compromise".*

[1] We are not including the welfare of third parties such as society or voters. These are also different across outcomes but, in principle, they are not internalized by players.

Proof. Suppose that there exist two sets $\tilde{S}_1 \subseteqq S_1$ and $\tilde{S}_2(\tilde{S}_1) \subseteqq S_2$ such that in a PBE of the game $a_1(s_1) = \rho$ and $a_2(s_2) = \rho$ with positive probability for all $s_1 \in \tilde{S}_1$ and $s_2 \in \tilde{S}_2(\tilde{S}_1)$.[2] Denote by $\overline{s}_1 = \max_{s_1 \in \tilde{S}_1}$ and $\overline{s}_2 = \max_{s_2 \in \tilde{S}_2(\tilde{S}_1)}$. According to this PBE, once agent 2 has observed $a_1 = \rho$, the following inequality must be satisfied:

$$\int_{s_1 \in \tilde{S}_1} \beta_2(s_1, s_2) dF_1(s_1 \,|\, s_1 \in \tilde{S}_1, s_2) \geqslant \int_{s_1 \in \tilde{S}_1 \cap s_1 < s_2} \alpha_2(s_1, s_2) dF_1(s_1 \,|\, s_1 \in \tilde{S}_1, s_2)$$
$$+ \int_{s_1 \in \tilde{S}_1 \cap s_1 > s_2} \gamma_2(s_1, s_2) dF_1(s_1 \,|\, s_1 \in \tilde{S}_1, s_2) \;\forall\, s_2 \in \tilde{S}_2(\tilde{S}_1)$$

where the l.h.s. is agent 2's expected payoff if $a_2 = \rho$ and the r.h.s. is his expected payoff if $a_2 = \phi$. This condition must hold in particular for $s_2 = \overline{s}_2$. Since $\alpha_2(s_1, s_2) > \beta_2(s_1, s_2) > \gamma_2(s_1, s_2)$, the inequality necessarily implies that $s_1 < \overline{s}_2$ must be binding at least for some $s_1 \in \tilde{S}_1$. Therefore, $\overline{s}_2 < \overline{s}_1$. Now, agent 1's decision is relevant only if $a_2 = \rho$. Thus, for the strategy described above to be a PBE, the following inequality must also hold:

$$\int_{s_2 \in \tilde{S}_2(\tilde{S}_1)} \beta_2(s_1, s_2) dF_2(s_2 \,|\, s_1) \geqslant \int_{s_2 \in \tilde{S}_2(\tilde{S}_1) \cap s_2 < s_1} \alpha_1(s_1, s_2) dF_2(s_2 \,|\, s_1)$$
$$+ \int_{s_2 \in \tilde{S}_2(\tilde{S}_1) \cap s_2 > s_1} \gamma_1(s_1, s_2) dF_2(s_2 \,|\, s_1) \;\; \forall\, s_1 \in \tilde{S}_1$$

Using the same reasoning as before, $\overline{s}_1 < \overline{s}_2$. Since both inequalities cannot be satisfied at the same time, $\tilde{S}_1 \neq \emptyset$ and $\tilde{S}_2(\tilde{S}_1) \neq \emptyset$ cannot both occur in equilibrium. □

The intuition is simple. In this class of games, agents know that information which is positive for them is negative for their rival. Thus, they have opposite interests on when to reach a compromise. As a result, whenever one agent wants to compromise, the other should not want to. For instance, country 1 has an incentive to stay in peaceful coexistence whenever its military strength $s_1$ is low. However, this is precisely when country 2 wants to force a war. In other words, in these games, one agent's gain is always the other agent's loss (of same or different magnitude, it does not matter). Since a compromise is broken as soon as one agent does not find it profitable, the fact that an agent wants to deal implies that the other should not accept

[2] Positive probability rather than probability 1 takes care of pure and mixed strategies at the same time.

it, and viceversa. The bottom line is that, in equilibrium, compromises are never possible. We want to stress the generality of this result, which holds for *any distribution* of strengths (the same or different for both players) and *any correlation* between the players' strengths. Since the results holds for any payoffs satisfying $\alpha_i > \beta_i > \gamma_i$, it means that introducing risk-aversion would not change the outcome of the game either. The result can be further extended as follows.

**Corollary 1** *The outcome of the game is still "no-compromise" if agents play the same stage game repeatedly and if agents announce their strategy simultaneously.*

Applying the same logic as before, repetition will not change the outcome of the game. Consider a two-stage version. If the second stage is reached, agents will update their beliefs about the type of their rival. However, since the no-compromise result does not depend on the functional form of the distribution of types, this revision of beliefs will not help reaching a compromise. Anticipating no-compromise in the second stage, parties will not compromise in the first stage either, for exactly the same reasons as in Proposition 1.

As for the simultaneous case, the only difference with the sequential game is that agent 2 will not compare his options conditional on having observed the choice of agent 1. However, it does not make any difference since, both in the sequential and the simultaneous versions, his action is only relevant if agent 1 offers a compromise. Thus the outcome of the Bayesian Nash Equilibrium (BNE) is, just like for the PBE, always no-compromise.

# 3 Laboratory experiment

## 3.1 Description of the game

This is a simplified version of the game described earlier. Each agent independently draws a number from a uniform distribution on $[0, 1]$ and privately observes their own number, which we refer to as the player's strength, $s_i$. Agent 1 chooses whether to "fight", $\phi$, or "retreat", $\rho$. If 1 chooses $\phi$, then the game ends. The agent with highest strength receives a win payoff $H$ and the other agent receives a lose payoff $L$ ($< H$). If agent 1 chooses $\rho$, then it is agent 2's turn. If agent 2 chooses $\phi$, then as before, the agent with highest strength receives a payoff of $H$ and the other receives a payoff of $L$. If, instead, agent 2 also chooses $\rho$, then agent 1 and agent 2 each obtains a pre-specified "compromise payoff" $M$, where $L < M < H$. Thus, the main simplification relative to the theoretical model presented in section 2 is that the win, lose and compromise payoffs are all independent of $(s_1, s_2)$. Each player's strength only affects payoffs vía the likelihood of winning under no-compromise. We look at several variations on this game.[3]

Variant 1. $H = 1$, $L = 0$, $M = .50$ with sequential move.
Variant 2. $H = 1$, $L = 0$, $M = .39$ with sequential move.
Variant 3. $H = 1$, $L = 0$, $M = .50$ with simultaneous move.
Variant 4. $H = 1$, $L = 0$, $M = .39$ with simultaneous move.

## 3.2 Relation to the experimental literature

We are not aware of any laboratory experiments of sequential games with two-sided private information. However, our setting shares several features with some well-known games. Below we describe (from most to least similar) two simultaneous games of multi-sided asymmetric information, two sequential games of one-sided asymmetric information and one static game of full information.

*The betting game* (Sonsino et al. (2001), Sovic (2004), Camerer et al. (2006)). An asset yielding a fixed surplus can be traded between agents who have private information. Trade occurs only if both agents agree. All the BNE of this game imply no trade. The information structure in this game is simpler than in ours: in 2 out of the 4 possible states, one agent has full information. The common knowledge of this information partition triggers unravelling to no-trade. This special partition is likely to facilitate learning. The authors do not study a sequential version of the game.

*Auction of a common value good and the winner's curse* (Kagel and Levin, 2002). As in our game, agents will play suboptimally if they do not anticipate the information contained in the rival's action. Our game allows for some simple comparative statics (different timings and different compromise payoffs). Also, our BNE and PBE are simple to compute. Again, we expect more rapid learning in the auction game. A bidder who does not realize the winner's curse and shades his bid accordingly is likely to win, lose money, realize his mistake, and learn for the next round.

[3]The nominal payoffs in the experiment are: $H = 95$, $L = 5$, $M \in \{50, 40\}$. We present here the scaled version $(x - 5)/90$.

*Adverse selection game* (Akerlof, 1970). The seller has a privately known valuation $\theta$. The buyer has a valuation *function* $B(\theta) > \theta$ but only knows the distribution from which $\theta$ is drawn. The buyer proposes a price and the seller accepts or refuses. As in our model, this game predicts some unravelling. However, the robust conclusion is the existence of a cutoff below which there is agreement or trade and above which there is not. This cutoff can be the lower bound (i.e., never agree as in our game), but it can also be the upper bound (i.e., always agree) or an interior value, depending on the parameters of the game. Samuelson and Bazerman (1985) show that the probability that buyers engage in unfavorable trades is increasing in the complexity of the adverse selection game. Note that because it is one-sided asymmetric information, the buyer's action has no signaling value. There have also been several market experiments with informed sellers and asymmetric information about product quality (Lynch et al., 1984).

*Blind bidding game* (Forsythe et al., 1989). This experiment asks a quite different question: will an informed seller reveal the quality of his good to the uninformed buyers? Full revelation occurs because the seller with the highest quality good has always an incentive to announce it, then so does the seller with second highest quality good, and so on. However, there is no role for the key effect of our game, namely the anticipation of information conveyed by the rival's action.

*Beauty contest* (Nagel, 1995). As our game, it predicts unravelling independently of the specific parameters of the problem. Since it is a static game of complete information, the reasons for convergence are different. Note also that even the most naïve learning rule ('play optimally given the outcome in the past round and assuming that nobody else revises his strategy') predicts rapid convergence if the game is played repeatedly. The experimental data confirms this prediction.

## 3.3 Experimental design and procedures

We conduced five sessions with a total of 56 subjects, using a simple $2 \times 2$ block design. The subjects were registered Princeton students who were recruited by email solicitation, and all sessions were conducted at The Princeton Laboratory for Experimental Social Science. All interaction in a session was computerized, using an extension of the open source software package, Multistage Games.[4] No subject participated in more than one session. The two dimensions of treatment variation were the compromise payoff ($M = .50$ vs. $M = .39$) and the order of moves (simultaneous vs. sequential play). In each session, subjects made decisions over 40 rounds, with $M$ fixed throughout the session. Half of the subjects participated in sessions with $M = .39$, and half the subjects participated in sessions with $M = .50$. In all sessions, we set $H = 1$ and $L = 0$. Each subject played exactly one game with one opponent in each round, with random rematching after each round. At the beginning of each round, $t$, each subject was independently assigned a new strength, $s_{it}$, drawn from a uniform distribution on $[0, 1]$.[5] Each subject observed his own strength, but had to make the fight-retreat decision before observing the strength of the subject they were matched with. The opponent's strength was revealed only at the end of the round.

At the beginning of each session, instructions were read by the experimenter standing on a stage in the front of the experiment room, which fully explained the rules, information structure, and client GUI for the *simultaneous move* game. A sample copy of the instructions is in the Appendix. After the instructions were finished, two practice rounds were conducted, for which subjects received no payment. After the practice rounds, there was an interactive computerized comprehension quiz that all subjects had to answer correctly before proceeding to the paid rounds. For the first 20 paid rounds of a session, subjects played the simultaneous version of the game. At the end of round 20, there was a brief instruction period during which rules for the sequential version of the game were explained.[6] In each match of the sequential version, one of the two players was randomly selected to be the first mover. After the first mover made a fight-retreat choice, the second mover was informed of that choice, but was not informed of the strength of the first mover. If the first mover's choice was fight, the second mover had no choice, and simply clicked a button on the screen labeled "continue". If the first mover's choice was retreat, the second mover had a choice between fight and retreat. After the second mover made a choice, the match ended

[4]Documentation and instructions for downloading the software can be found at http://multistage.ssel.caltech.edu.

[5]In the experimental implementation of payoffs, the $H$ and $L$ payoffs paid off \$.57 and \$.03, respectively. The compromise payoff $M$ was scaled accordingly, at \$.30 and \$.24 for the two treatments.

[6]In one of the sessions, the sequential version was played in rounds $1 - 20$ and the simultaneous version was played in rounds $21 - 40$.

and the strength levels and outcome were revealed. The subjects then participated in 20 additional rounds of the sequential version of the game, with opponents, roles (first or second mover), and strengths randomly reassigned at the beginning of each round. Subjects were paid the sum of their earnings over all 40 paid rounds, in cash, in private, immediately following the session. Sessions averaged one hour in length, and subject earnings averaged $25. Table 1 displays the pertinent details of the five sessions.

| Session | # subjects | $M$ | rounds 1-20 | rounds 21-40 |
|---|---|---|---|---|
| 1 | 8 | .50 | sequential | simultaneous |
| 2 | 8 | .50 | simultaneous | sequential |
| 3 | 12 | .50 | simultaneous | sequential |
| 4 | 14 | .39 | simultaneous | sequential |
| 5 | 14 | .39 | simultaneous | sequential |

**Table 1**. Session details for the experiment.

# 4 A descriptive analysis of the results

In this section, we provide a descriptive analysis of the experimental results. We discuss the main aggregate features of the data, including the mean rates of fight and retreat, both overall and as a function of strength, and explore time trends. We compare the data to two natural benchmarks. The first benchmark is Nash equilibrium, in which all players always choose $\phi$ regardless of strength (and, in the sequential version, regardless of the choice of the first mover). A second, weaker benchmark is the type-independent model, where the probability of fighting is independent of strength. We study the differences in probabilities of fighting as a function of the compromise payoff and the timing of the game. Last, we analyze the data at an individual level. For each player, we estimate a decision rule that maps strength into a probability of fighting.

## 4.1 Aggregate fight rates unconditional on strength

The simplest cut at the data is to compare the relative frequencies of choosing $\rho$ or $\phi$, without conditioning on the actual draws of $s_i$. Table 2 shows the relative frequencies of $\phi$ in the experiment, broken down by compromise payoff and order of moves. The number of subjects is in parenthesis.

| Order | Role | $M = .39$ | $M = .50$ |
|---|---|---|---|
| Sequential | First | .589 (280) | .538 (264) |
| Sequential | Second | .643 (115) | .566 (122) |
| Simultaneous | — | .657 (560) | .573 (560) |

**Table 2**. Unconditional frequency of choosing $\phi$.

There are several interesting comparisons. First, for all roles, there is a difference in the simultaneous treatment and for both roles in the sequential treatment between $\phi$ rates in the $M = .39$ and the $M = .50$ treatments. Fighting is chosen less frequently when the "compromise dividend" is higher. The differences are fairly small in magnitude and significant at the 5% level only for the simultaneous treatment.[7] Second, there is a small, but not significant difference between the $\phi$ rates for the first and second movers in both the $M = .39$ and the $M = .50$ treatments: first movers in the sequential game choose $\phi$ less frequently than second movers and also less frequently than players of the simultaneous game. Third, second movers also choose $\phi$ somewhat less frequently than players in the simultaneous games, in both the $M = .39$ and the $M = .50$ treatments, but the differences are insignificant.

The logic of the game suggests that, over time, learning will lead to unravelling. That is, perceptive players should be able to realize that they will improve their payoff by adopting a cutoff strategy lower than the cutoff strategy used by their opponent. Given the symmetry of the game, they should realize that perceptive opponents will also notice this. The unravelling logic may be responsible for the higher fighting rates of second relative to first movers. It also suggests that $\phi$ rates should be increasing over time, in all treatments. We investigate this hypothesis by breaking the data down into early and late matches. In each session, there were 20 rounds each of the sequential and the simultaneous games. We code the choices in the first 10 rounds of each version of the game as "inexperienced" and the last 10 rounds of each version as "experienced". Table 3 presents the $\phi$ rates, broken down by experience level. The number of subjects is in parenthesis.

[7]The difference is significant at the 1% level in the sequential treatment if the two roles are pooled.

| Order | Role | $M = .39$ | | $M = .50$ | |
|---|---|---|---|---|---|
| | | inexper. | exper. | inexper. | exper. |
| Sequential | First | .564 (140) | .614 (140) | .484 (124) | .586 (140) |
| Sequential | Second | .672 (61) | .611 (54) | .484 (64) | .655 (58) |
| Simultaneous | — | .611 (280) | .704 (280) | .582 (280) | .564 (280) |

**Table 3**. Unconditional frequency of choosing $\phi$ by experience level.

The effects of experience on the unconditional $\phi$ rates is ambiguous. In four of the six comparisons, the $\phi$ rate increases, as hypothesized, although it remains well below 1. All four such differences are statistically significant. In two of the six comparisons, $\phi$ decreases, but these two changes are not significant. Furthermore, the two treatments where $\phi$ decreases have no apparent relation with each other (simultaneous with $M = .50$ and second player in sequential with $M = .39$). Some possible explanations for this very weak of evidence for learning are discussed later in the paper.

## 4.2 Aggregate fight rates conditional on strength

The analysis above, while providing a useful sketch of the results, falls short of giving a complete picture of the aggregate data, because the unconditional $\phi$ rates are not a sufficient statistic for the actual strategies. A behavior strategy in each game is a probability of choosing $\phi$ *conditional on* $s$. By aggregating across all the (strength, action) paired observations for a treatment, we can graphically display the aggregate empirical behavior strategy, and then compare this strategy across treatments. Figure 1 shows six graphs. The graphs on the left correspond to $M = .39$, and the graphs on the right are for $M = .50$. The middle and bottom graphs are for the first and second movers in the sequential treatment, and the top graphs for the simultaneous movers. The strength is on the horizontal axis, on a scale of 0 to 100, and the empirical fighting frequencies are on the vertical axis on a scale of 0 to 1. Thus, for example, if all subjects were to choose the same cutoff strategy $s^*$, then we would observe a step function, with a probability of fighting equal to 0 below $s^*$ and equal to 1 above $s^*$. Note that functions need not be monotonically increasing, although we expect that players with higher strength will be more likely to fight. The empirical fighting frequencies are moving averages over 5 strength levels.

[ Insert Figure 1 ]

These graphs suggest that the second-movers in the sequential version of the game behave differently in at least two ways. First, they tend to fight more. If one looks at the point in the graph where the fight probabilities first reach 50%, this switchpoint is in the high .20s for second movers in both the .39 and .50 treatments, while it is in the mid to high .30s for simultaneous movers and even higher for the first movers in the sequential treatment.

The second movers also display less erratic behavior, in the sense that for low values they (almost) never fight and for high values they (almost) always fight. This is reflected in a steeper response curve.

This latter observation begs for an explanation. Since the decision of a player matters only if the rival chooses $\rho$, players should condition their action on that event, rendering irrelevant their position in the game (first, second or simultaneous mover). The increased fighting and less erratic behavior of second movers suggests instead that players find it easier to respond to *observed* than to *anticipated* choices of the rival. This relates to the findings on the winner's curse where bidders have a difficult time in conditioning the bid on the information conveyed by their winning the auction.

## 4.3 Individual cutpoint analysis

In order to address the question of conditional fight rates more carefully, we turn to an analysis of individual choice behavior. Do subjects adopt cutpoint decision rules? If so, how much variation is there across individuals? How do cutpoints vary across treatments? We document that indeed nearly all subjects use (approximate) cutpoint strategies, there is some heterogeneity across subjects, and the distribution of these strategies varies systematically across treatments.

In order to estimate decision rules, we use a simple optimal classification procedure, similar to Casella et al. (2006) and Palfrey and Prisbrey (1996). For each subject and each condition the subject is in, we look at the set of strengths they were randomly assigned, and the corresponding fight/retreat decision they made. For any hypothetical cutpoint strategy for an individual subject, we can then ask how many of these decisions are correctly classified. For example, if in some round a subject with strength 40 chose $\phi$, the decision would be correctly classified only if the hypothetical cutpoint were less than or equal to 40. We then use the hypothetical cutpoint with the fewest misclassified decisions as the estimate for that individual/condition.[8] Table 4 reports the average estimated cutpoint across all subjects, and the percentage of misclassified decisions, by condition. The average cutpoints mirror the aggregate fight rates by condition reported in Table 2. The most fighting occurs in .39 treatments, and there is more fighting (and fewer "errors") by second movers than first movers.

| Condition | $M$ | Average Estimated Cutpoint | % misclassified |
|---|---|---|---|
| Simultaneous | .39 | 36.7 | 3.4 |
| Simultaneous | .50 | 45.5 | 3.7 |
| First Mover | .39 | 40.2 | 2.1 |
| First Mover | .50 | 45.1 | 2.3 |
| Second Mover | .39 | 37.0 | 0.0 |
| Second Mover | .50 | 40.7 | 0.8 |

**Table 4**. Cutpoint summary statistics.

Overall, very few decisions are misclassified. In each of the simultaneous treatments, 16 out of 28 subjects are perfectly classified. In the sequential conditions, the number of perfectly classified subjects range from 23 to 28 out of 28. The worst case of misclassification was one subject in the simultaneous, .50 treatment who has 5 misclassified observations. Our inference from this is that subjects use cutpoint strategies, with rare exceptions.

The next question is whether subjects are using the same cutpoints, or if instead there is a significant amount of heterogeneity. Indeed, we find quite a bit of heterogeneity. Figure 2 displays the cumulative frequency distribution of estimated cutpoints for all the treatments. The horizontal axis represents hypothetical cutpoints, ranging from 0 to 100. The vertical axis indicates how many subjects (out of 28) were estimated to have been using a cutpoint less than or equal to that number.

[ Insert Figure 2 ]

So, for example, in the simultaneous, $M = .39$ treatment, the distribution of estimated cutpoints is approximately Uniform between 20 and 50. Graphs of the distribution of cutpoints in other conditions also exhibit a wide range of estimated cutpoints, with few above 60 or below 20.

[8] If there are multiple best fitting cutpoints, we take the average.

## 4.4 Summary of descriptive analysis

The main findings of our analysis so far can be summarized as follows. First, the unconditional $\phi$ rates range from about 50% to 70%, depending on the treatment, falling far short of the theoretical prediction of 100%. Second, in all treatments, the $\phi$ rate conditional on strength increases monotonically from virtually 0% for strengths below 20% to virtually 100% for strengths above 60% if $M = .39$ or above 70% if $M = .50$. Third, comparative statics with respect to $M$ has the expected sign, with less fighting when the compromise payoff is higher. Fourth, second movers display more fighting and less erratic behavior than first and simultaneous movers. This suggests the importance of actually observing the behavior of the rival before making inferences and choices, rather than just conditioning on a hypothetical event. Fifth, there is little evidence of learning: unconditional fight rates increase significantly over time in 4 out of 6 conditions but decrease in the other 2 conditions. A possible explanation is the insufficient feedback provided to players (only the rival's strength and the outcome is revealed at the end of each round). However, given that the order of moves matters, one would expect that second movers would use their knowledge of unravelling when they subsequently play first.[9] Sixth, the vast majority of subjects use cutpoint strategies. This shows their understanding, at least at an intuitive level, that the expected payoff differential between $\phi$ and $\rho$ increases with the player's strength. Seventh, there is substantial heterogeneity in the players' cutpoints, and the distribution of these cutpoints varies by condition in ways that mirror the differences in the aggregate fight rates.

# 5 Competing models to explain the data

In this section, we consider several models to explain the excessively low fight rates. Note that this game is easily solved by iterated dominance, but only using weak rather than strict dominance. Denote a strategy as a function $q : [0,1] \to [0,1]$. First note that any strategy $q$ that assigns $q(1) < 1$ is weakly dominated by the strategy $q'$ where $q'(s) = q(s)$ for all $s \neq 1$ and $q'(1) = 1$. In the experiments, the type distribution was discrete, so once we

[9] Recall that in our design, all players gain experience as *both* first and second movers. That is, our data on first and second movers are all coming from the *same* subjects. Subjects apparently do not draw inferences from their own decision making in different roles about how other subjects behave in those roles.

eliminate all those strategies, then any strategy $q$ that assigns $q(.99) < 1$ and $q(1) = 1$ is weakly dominated by the strategy $q'$ where $q'(s) = q(s)$ for all $s \neq .99$ and $q'(.99) = 1$. And so forth.[10] On the other hand, rationalizability does not eliminate any strategy, since *every* strategy is a weak best response to the equilibrium strategy, $q^*(s) = 1$ for all $s$.

The three models we consider all have features that admit the possibility of observing weakly dominated strategies. The first such model is quantal response equilibrium or QRE (McKelvey and Palfrey, 1995); the second is cognitive hierarchy or CH (Camerer et al., 2004); and the third is cursed equilibrium or CE (Eyster and Rabin, 2005).[11] We also consider some variations of these models that allow for heterogeneity or hybridize between models, such as the truncated quantal response equilibrium or TQRE (Camerer et al., 2006).

## 5.1 Quantal Response Equilibrium

Quantal response equilibrium applies stochastic choice theory to strategic games, and is motivated by the idea that a decision maker may take a suboptimal action, and the probability of doing so is increasing in the expected payoff of the action. In a regular QRE (Goeree et al., 2005), one simply replaces the best response function used to characterize Nash equilibrium, with a quantal response function that is continuous and monotone in expected payoffs. That is, the probability of choosing a strategy is a continuous increasing function of the expected payoff of using that strategy, and strategies with higher payoffs are used with higher probability than strategies with lower payoffs. A quantal response equilibrium is then a fixed point of the quantal response mapping. In a logit equilibrium, for any two strategies, the log odds of the choice probabilities are proportional to the difference in expected payoffs, where the proportionality factor, $\lambda$, is a measure of responsiveness of choices to payoffs. That is:

$$\ln\left[\frac{\sigma_{ij}}{\sigma_{ik}}\right] = \lambda\Big[U_{ij} - U_{ik}\Big]$$

[10]For the $M = .50$ game, at the last iteration, a player with the lowest strength, $s = .01$, is indifferent between $\phi$ and $\rho$, and therefore, there is an equilibrium with $s = .01$ types choosing $\rho$ and all other types choosing $\phi$. For the $M = .39$ game, the iteration continues all the way down, and the only equilibrium is $q^*(s) = 1$ for all $s$.

[11]Some preliminary findings about CH and CE are discussed in Wang (2006), with permission of the authors.

where $\sigma_{ij}$ is the probability $i$ chooses strategy $j$ and $U_{ij}$ is the corresponding expected payoff in equilibrium. Note that a higher $\lambda$ reflects a "more precise" response to the payoff differential. The polar cases $\lambda = 0$ and $\lambda \to +\infty$ correspond to the cases of pure noise and Nash equilibrium, respectively.

Given this stochastic choice model, it is natural to first ask how costly are suboptimal strategies in the compromise game. Obviously, there is no absolute answer to this question, since there are many ways to play non-optimally, but the following analysis indicates that the expected costs of suboptimal behavior are quite small.

**Expected losses from suboptimal cutpoint strategies** Suppose that agent 1 plays a (suboptimal) cutpoint strategy $\tilde{s}_1 \in (0,1)$ and player 2 best responds to that strategy. If $a_1 = \rho$, agent 2's expected payoff of playing $a_2 = \rho$ and $a_2 = \phi$ are, respectively, $M$ and $\Pr[s_2 > s_1 \,|\, s_1 < \tilde{s}_1]\, H + \Pr[s_2 < s_1 \,|\, s_1 < \tilde{s}_1]\, L$. One can immediately see that there is a cutoff $\tilde{s}_2 < \tilde{s}_1$ such that $a_2(s_2) = \phi$ if $s_2 \geq \tilde{s}_2$ and $a_2(s_2) = \rho$ if $s_2 < \tilde{s}_2$. Under a uniform distribution, $H = 1$ and $L = 0$, we have $\tilde{s}_2 = M\tilde{s}_1$. Let $V_i(\tilde{s}_1, \tilde{s}_2)$ be agent $i$'s expected payoff given the cutpoint strategies $\tilde{s}_1$ and $\tilde{s}_2 = M\tilde{s}_1$. We have:

$$V_2 = M \int_0^{\tilde{s}_2} \int_0^{\tilde{s}_1} ds_1\, ds_2 + \int_{\tilde{s}_2}^1 \int_0^{s_2} ds_1\, ds_2 = \frac{1}{2} + \frac{M^2 \tilde{s}_1^2}{2}$$

$$V_1 = M \int_0^{\tilde{s}_1} \int_0^{\tilde{s}_2} ds_2 ds_1 + \int_{\tilde{s}_2}^{\tilde{s}_1} \int_{\tilde{s}_2}^{s_1} ds_2 ds_1 + \int_{\tilde{s}_1}^1 \int_0^{s_1} ds_2 ds_1 = \frac{1}{2} - \frac{M(2-3M)\tilde{s}_1^2}{2}$$

In the Nash equilibrium, $\tilde{s}_1 = \tilde{s}_2 = 0$ so $V_i = 1/2$: both players win with equal ex-ante probability. If, instead, agent 1 plays a naïve strategy $\tilde{s}_1 = 1/2$, his ex-ante loss is about 6% of his expected payoff if $M = .50$ and 8% if $M = .39$, two relatively small numbers. Also, the computations are made under the assumption that player 2 perfectly anticipates player 1's strategy. Therefore, these numbers provide an upper bound in the expected cost of playing suboptimally. Last, it is interesting to notice that if the compromise payoff is sufficiently high ($M \geq .67$), then $V_i(\tilde{s}_1, M\tilde{s}_1) > V_i(0,0)$ for all $i, \tilde{s}_1 > 0$. That is, *both* agents benefit if one of them does not play the optimal strategy even if this deviation is anticipated by the rival.

Overall, the analysis suggests that the expected cost of using a suboptimal cutpoint strategy is likely to be small. This is confirmed by the data. Using the empirical choice frequencies in the experiment, we calculated the expected payoffs from different cutpoint strategies. Indeed these are rather flat, as illustrated in Figure 3, which graphs these payoffs for the simultaneous, $M = .50$ treatment (the other treatments have similar results and thus are not reported here). This result may partly explain why there is so little evidence of learning in the experiment.

[ Insert Figure 3 ]

**Specification of the QRE model** We consider two different specifications of the logit equilibrium version of QRE. The first specification takes an interim approach and analyzes the game in behavioral strategies. This approach corresponds to the agent QRE (AQRE) of McKelvey and Palfrey (1998). Conditional on player 1's strength, and given the AQRE behavioral strategies used by player 2, the log-odds of player 1 choosing retreat vs. fight is proportional to the difference in expected payoffs between retreat and fight – and similarly for player 2.

The second analyzes the game in ex ante strategies, and assumes players choose stochastically over possible plans for whether or not to fight as a *function* of strength.[12] Because the set of all possible pure strategies in our game is huge ($2^{100}$), we are forced to consider only a subset of such strategies. The natural restriction is to consider only monotone strategies, i.e., cutpoint strategies. This is a natural restriction since monotone strategies are always best responses and furthermore any non-monotone strategy is weakly dominated by a monotone strategy. This also reduces the set of pure strategies to a small enough number (100) that estimation is possible. The above analysis about the expected cost of using suboptimal cutpoints is especially relevant for this cutpoint QRE approach.

In the logit parameterization of the cutpoint QRE, the distribution over cutpoint strategies used by player 2 has the standard property. Namely, the log-odds of player 1 choosing any cutpoint $c$ versus any other cutpoint $c'$ is proportional to the *ex ante* difference in expected payoffs between using those two cutpoints – and similarly for player 2.

**Logit QRE in behavioral strategies** For any response parameter $\lambda$ we solve for a fixed point in behavioral strategies. Denote by $\phi^*_\lambda$ such an equilibrium fixed point, and by $\phi^*_\lambda(s)$ the equilibrium probability of $\phi$ given $s$.

[12] A similar approach is taken in the ex ante QRE model explored in Casella et al. (2006).

Consider first the simultaneous game. We need to determine the expected utility of $\phi$ for a player with strength $s$ conditional on the other player using strategy $\phi^*_\lambda$ and having chosen $\rho$. This is simply equal to the conditional probability that the other player has strength less than $s$, given that he has chosen $\rho$. It is then given by:

$$V_\phi(s;\phi^*_\lambda) = \frac{\int_0^s [1-\phi^*_\lambda(t)]dt}{\int_0^1 [1-\phi^*_\lambda(t)]dt} \tag{1}$$

The expected utility of $\rho$ conditional on the other player having chosen $\rho$ is simply $V_\rho(s;\phi^*_\lambda) = M$, so the difference in the expected utility of $\phi$ and $\rho$ is:

$$\begin{aligned}\Delta(s;\phi^*_\lambda) &= \int_0^1 [1-\phi^*_\lambda(t)]dt \left(V_\phi(s;\phi^*_\lambda) - V_\rho(s;\phi^*_\lambda)\right) \\ &= \int_0^s [1-\phi^*_\lambda(t)]dt - M\int_0^1 [1-\phi^*_\lambda(t)]dt\end{aligned}$$

Hence, in a symmetric logit QRE, $\phi^*_\lambda$ is characterized by:

$$\phi^*_\lambda(s) = \frac{e^{\lambda\Delta(s;\phi^*_\lambda)}}{1+e^{\lambda\Delta(s;\phi^*_\lambda)}} \quad \text{for all } s \in [0,1]$$

The sequential game requires solving simultaneously for $\phi^*_{\lambda 1}(s_1)$ and $\phi^*_{\lambda 2}(s_2)$. The expressions for the first mover are exactly the same as in the simultaneous move game, so, modifying the notation slightly to make clear that it is player 1's equation, we get:

$$\phi^*_{\lambda 1}(s_1) = \frac{e^{\lambda\Delta_1(s_1;\phi^*_{\lambda 2})}}{1+e^{\lambda\Delta_1(s_1;\phi^*_{\lambda 2})}} \quad \text{for all } s_1 \in [0,1]$$

where

$$\Delta_1(s_1;\phi^*_{\lambda 2}) = \int_0^{s_1} [1-\phi^*_{\lambda 2}(t)]dt - M\int_0^1 [1-\phi^*_{\lambda 2}(t)]dt$$

The condition for the second mover is the same, except the second mover's expected utility difference does not have to be conditioned on the first mover choosing $\rho$, so we get:

$$\phi^*_{\lambda 2}(s_2) = \frac{e^{\lambda\Delta_2(s_2;\phi^*_{\lambda 1})}}{1+e^{\lambda\Delta_2(s_2;\phi^*_{\lambda 1})}} \quad \text{for all } s_2 \in [0,1]$$

where

$$\Delta_2(s_2;\phi^*_{\lambda 1}) = \frac{\int_0^{s_2} [1-\phi^*_{\lambda 1}(t)]dt}{\int_0^1 [1-\phi^*_{\lambda 1}(t)]dt} - M$$

**Logit QRE in cutpoint strategies** We next consider the slightly more sophisticated version of QRE where players are assumed to randomize over monotone cutpoint strategies, which we call QRE-cut. In our game, a cutpoint strategy is a critical value of strength, $c$, such that player $i$ chooses $\phi$ if $s_i \geq c$ and chooses $\rho$ if $s_i < c$. Hence, we define a *cutpoint* quantal response to be given by two probability distributions over $c$, one for each player, denoted $q_1(c)$ and $q_2(c)$. In the simultaneous version of the game, we consider only symmetric QRE-cut where $q_1(c) = q_2(c) = q(c)$ for all $c$. For the sequential version, generally $q_1(c) \neq q_2(c)$ since it is not a symmetric game, and the second player chooses a cutpoint after observing the first player's move. We use the logit quantal response function for a parametric specification. Hence, the probability that a player chooses a particular strategy is proportional to the exponentiated expected payoff from using that strategy, given the cutpoint quantal response function of the other player. It is worth noting that past studies have found that in binary choice games with continuous types, a cutpoint strategy can be a useful variation on the standard QRE approach (see Casella et al., 2006). Furthermore, the analysis in section 4.3 suggests that subjects seem to adhere to this type of strategies.

Consider the simultaneous game. The expected utility to player 1 of using a cutpoint strategy $\tilde{c}$ if player 2 uses $q(\cdot)$ is given by:

$$U(\tilde{c}) = \int_{\tilde{c}}^{1} s ds + \int_{0}^{\tilde{c}} \left[ \int_{0}^{s} q(c) \left(cM + (s - c)\right) dc + \int_{s}^{1} q(c) cM dc \right] ds \qquad (2)$$

The first term is the probability of drawing a strength $s$ above the cutpoint, in which case player 1 chooses $\phi$ and obtains a payoff 1 only if player 2 has a lower strength. The second term is the probability of drawing a strength $s$ below the cutpoint, in which case player 1 chooses $\rho$. Then, if player 2's strength is lower, a compromise gives payoff $M$ and a no-compromise gives payoff 1, and if player 2's strength is higher then a compromise gives payoff $M$ and a no-compromise gives payoff 0. In a symmetric logit QRE-cut:

$$q(\tilde{c}) = \frac{e^{\lambda U(\tilde{c})}}{\int_0^1 e^{\lambda U(c)} dc} \quad \text{for all } \tilde{c} \in [0, 1]$$

In the sequential game, the expression for the first mover's utility of using $\tilde{c}$, given player 2 uses $q_2(\cdot)$ is the same as in the simultaneous case:

$$U_1(\tilde{c}) = \int_{\tilde{c}}^{1} s_1 ds_1 + \int_{0}^{\tilde{c}} \left[ \int_{0}^{s_1} q_2(c) \left(cM + (s_1 - c)\right) dc + \int_{s_1}^{1} q_2(c) cM dc \right] ds_1 \qquad (3)$$

By contrast, the second mover's utility of using $\tilde{c}$, given player 1 uses $q_1(\cdot)$ does not have to be conditioned on the first mover choosing $\rho$. That is:

$$U_2(\tilde{c}) = \int_{\tilde{c}}^{1} \left[ \frac{\int_0^{s_2} c_1 q_1(c_1) dc_1}{\int_0^1 c_1 q_1(c_1) dc_1} + \frac{\int_{s_2}^{1} s_2 q_1(c_1) dc_1}{\int_0^1 c_1 q_1(c_1) dc_1} \right] ds_2 + \tilde{c} M \quad (4)$$

There are three observations to make about the cutpoint QRE solutions. First, in the sequential game, the equilibrium cutpoint distributions are different for the two players. The second mover generally adopts lower cutpoints, which translates into higher $\phi$ rates in the figures. Second, players adopt lower cutpoints when $M$ is lower. Third, the cutpoint distributions for the first mover in the sequential games are different from the cutpoint distributions in the corresponding simultaneous games, even though the utility formulas (equations 2 and 3) are identical.

We fit the behavioral strategy logit QRE and the cutpoint strategy QRE models by standard maximum likelihood techniques, i.e., finding the value of $\lambda$ that maximizes likelihood of the observed frequencies of strategies. We estimated restricted and unrestricted versions of the models. In the most restricted version, the parameters are constrained to be the same across all treatments. We also estimate a version of the model where the parameters are constrained to be the same for the .39 and .50 treatments, but are allowed to be different in the simultaneous and sequential games. The results for the cutpoint QRE are displayed in Table 5.[13]

## 5.2 Cognitive Hierarchy

The CH model (Camerer et al., 2004) postulates that when a player makes a choice, his decision process corresponds to a "level of sophistication" $k$ with probability $p_k$. CH models are completely determined by an assumption about how level 0 types behave ($\sigma_0$), and the distribution of levels of sophistication ($p$). Once the behavior of level 0 players is determined, level 1 players are characterized by choosing with equal probability all strategies that are best responses to level 0 opponents. Level 2 players optimize assuming they face a distribution of level 0 and level 1 players, where the distribution satisfies *truncated rational expectations*. That is, the beliefs of level 2 players that

[13]The estimates for the behavioral strategy version of QRE are not reported because the fit was significantly worse than for the cutpoint QRE in all treatments. This is not surprising given that most subjects were found to use cutpoints.

their opponent is choosing according to a level 0 or a level 1 decision process, denoted $b^2(0)$ and $b^2(1)$, is given by the truncated "true" distribution of these types: $b^2(0) = \frac{p_0}{p_0+p_1}$ and $b^2(1) = \frac{p_1}{p_0+p_1}$. Level 2 players are characterized by choosing with equal probability all strategies that are best responses to $b^2$ beliefs about the opponents. Higher levels are defined analogously, so a level $k$ optimizes with respect to beliefs $b^k$ where $b^k(j) = p_j / \sum_{l=0}^{k-1} p_l$ (see also Costa-Gomes and Crawford (2006) and Crawford and Iriberri (2005) for related models of hierarchical thinking).

For any distribution of levels, $p$, this implies a unique specification of a mixed strategy for each level, $\sigma(p) = (\sigma_0(p), ..., \sigma_k(p), ...)$, and this specification can be solved recursively, starting with the lowest types. This generates predictions about the aggregate distribution of actions, denoted $\overline{\sigma}(p) = \sum_{k=0}^{\infty} p_k \sigma_k(p)$. In all applications to date, $p$ is assumed to be Poisson distributed with mean $\tau$. That is, $p_k = \frac{\tau^k}{k!} e^{-\tau}$. We consider two specifications of the behavior of level 0 types.

**Random actions** In the standard CH model, level 0 players are typically assumed to choose an action randomly. In the context of our game, this means that they are equally likely to select $\phi$ or $\rho$, independently of their strength. Level 1 types best respond to level 0 types. It can be easily shown that the best response strategy is to choose cutpoint $M$. Level 2 players then optimize with a cutpoint somewhere between $M$ (the best response if everyone is level 0) and $M^2$ (the best response if everyone is level 1), with the exact value depending on $p_0$ and $p_1$. Behavior by higher level players is defined recursively.

**Random cutpoints** An alternative version, which we call the cutpoint CH model, simply replaces the assumption that level 0 types randomize uniformly over actions, with the assumption that they randomize uniformly over cutpoint strategies. This implicitly endows level 0 types with some amount of rationality, in the form of monotone behavior: they are more likely to choose $\phi$ when their strength is high than when their strength is low. It is immediate to see that, in our game, a level 0 type who randomizes over cutpoints has a probability of $\phi$ as a function of $s$ which is equal to $s$. As in the standard CH, the best responses of higher types will be unique cutpoints, and are easily calculated by recursion. Since a level 0 type has a probability $1 - s$ of choosing $\rho$, the posterior distribution of strength of a level 0 type conditional on choosing $\rho$ is $f(s \mid \rho) = \frac{1-s}{\int_0^1 (1-x)dx} = 2-2s$. Hence, the expected payoff of $\phi$ for a level 1 type with strength $s$ and conditional on the other player being level 0 and choosing $\rho$ is $\int_0^s (2-2x)dx = 2s - s^2$. Since the payoff of $\rho$ is $M$, the optimal cutpoint of a level 1 type is the value $s_1^M$ that solves $2s_1^M - (s_1^M)^2 = M$, that is $s_1^M = 1 - \sqrt{1-M}$. For our two treatments, we get $s_1^{.50} = 1 - \sqrt{1/2} \approx .29$ and $s_1^{.39} = 1 - \sqrt{11/18} \approx .22$. Higher types are then defined recursively, with the exact cutpoint for a level $k$ depending on $\{p_l\}_{l=0}^{k-1}$. This produces a CH model that is comparable to QRE in the sense that *all* players choose cutpoint strategies, so $\phi$ probabilities are monotone in $s$ for all players.

We fit the Poisson specification of the CH and cutpoint CH models to the dataset by finding the value of $\tau$ that maximizes likelihood of the observed aggregate frequencies of strategies, under the assumption that types are identically and independently distributed draws. We estimated the best-fitting values of $\tau$ by maximum likelihood for each of the four treatments, and present the results in Table 5. We report both constrained and unconstrained estimates.[14]

## 5.3 Combining quantal response and strategic hierarchies (TQRE)

The predictions of the CH model (both with random actions and random cutpoints) differ from the QRE and QRE-cut models in two important ways. First, in CH, all players with the same level of sophistication choose the same cutpoint strategy. Second, predictions in CH are identical for the sequential and simultaneous versions of the game.[15] Neither "bunching" by layers of reasoning nor identical behavior in the simultaneous and sequential treatments are observed in the data.

An approach that combines QRE and hierarchical thinking, called *Truncated Quantal Response Equilibrium (TQRE)*, is developed in Camerer et al.

[14]The random action version of CH did not fit nearly as well as the cutpoint CH model, so we do not report it here. Again, this is not surprising given that most subjects were found to use cutpoints.

[15]If we assume that level 0 types are not random players but just confused by the difficulty of the game, then it could be argued that a level 0 second mover can infer some information about how to play only by observing the action of the first mover. This would imply different choices for first and second movers in CH. However, the method to obtain those differences seems somewhat ad-hoc. We therefore decided not to explore this route.

(2006). This model introduces a countable number of players' *skill levels*, $\lambda_0, \lambda_1, ..., \lambda_k, ...$ . The distribution of skill levels in the population is given by $p_0, p_1, ..., p_k, ...$ . A player with skill level $k$ chooses stochastically with a logit quantal response function with precision $\lambda_k$. TQRE assumes truncated rational expectations in a similar manner to CH, so that a player with precision $\lambda_k$ has beliefs $p_j^k = p_j / \sum_{l=0}^{k-1} p_l$ for $j < k$ and $p_j^k = 0$ for $j \geq k$. For reasons of parsimony and comparability to CH, we assume that skill levels are Poisson distributed and equally space $\lambda_k = \gamma k$. Thus, it is a two parameter model with Poisson parameter, $\tau$ , and a spacing parameter, $\gamma$.

The TQRE model has two effects. It smoothes out the mass points, and it makes different predictions for the sequential and simultaneous games. These effects work slightly differently with behavioral strategies and with cutpoint strategies. As was true with QRE and CH separately, the cutpoint model does better, so we report only those estimates in Table 5.

## 5.4 Cursed Equilibrium

In a CE model, players are assumed to systematically underestimate the correlation between the opponents' action and information. As in the CH model, a cursed equilibrium will be the same in both the sequential and simultaneous treatments. In an $\alpha$-cursed equilibrium ($\text{CE}_\alpha$) all players are $\alpha$-cursed. However, players believe that opponents are $\alpha$-cursed with probability $(1 - \alpha)$ and they believe that actions of opponents are independent of their information with probability $\alpha$. All players optimize with respect to this (incorrect) mutually held belief about the joint distribution of opponents' actions and information. In our model, we can easily compute the cutpoint strategy in $\text{CE}_\alpha$ as a function of $M$, denoted $s_\alpha^*(M)$. For a player with strength $s_i$, and assuming the other player is using $s_\alpha^*(M)$, the expected utility of $\phi$, conditional on the opponent choosing $\rho$ is given by:

$$\begin{aligned} V_\phi^\alpha(s_i) &= \alpha \Pr\{\mathrm{s}_j < s_i\} + (1-\alpha)\Pr\{\mathrm{s}_j < s_i \,|\, a_j = \rho, s_\alpha^*(M)\} \\ &= \alpha s_i + (1-\alpha)\min\{1, \tfrac{s_i}{s_\alpha^*(M)}\} \end{aligned}$$

A player with strength equal to the equilibrium cutpoint must be indifferent between $\phi$ and $\rho$. Formally, $V_\phi^\alpha(s_\alpha^*(M)) = V_\rho^\alpha(s_\alpha^*(M))$. Therefore:[16]

$$s_\alpha^*(M) = \begin{cases} 1 - \frac{1-M}{\alpha} & \text{if} \quad \alpha > 1-M \\ 0 & \text{if} \quad \alpha \leq 1-M \end{cases}$$

A difficulty with the $\mathrm{CE}_\alpha$ model is that it cannot be fit to the data due to a zero-likelihood problem: for each value of $\alpha$ it makes a point prediction. Therefore, we slightly modify the equilibrium concept in order to allow for stochastic choice. The approach we follow is to combine QRE with $\mathrm{CE}_\alpha$.[17] In the simultaneous move game, a (symmetric) $\alpha$-QRE is a behavior strategy, or a set of probabilities of choosing $\phi$, one for each value of $s \in [0, 1]$. We denote such a strategy evaluated at a specific strength value by $\phi(s)$. Given $\lambda$ and $\alpha$ we denote $\alpha$-QRE as the behavior strategy $\phi_{\lambda\alpha}^*$. If player $j$ is using $\phi_{\lambda\alpha}^*$ and player $i$ is $\alpha$-cursed, then $i$'s expected payoff from choosing $\phi$ when $s_i = s$ is given by:

$$\begin{aligned} V_\phi^\alpha(s) &= \int_0^1 \phi_{\lambda\alpha}^*(t)dt\Big[\alpha s + (1-\alpha)\Pr\{s_j < s \,|\, a_j = \phi, \phi_{\lambda\alpha}^*\}\Big] \\ &\quad + \int_0^1 [1-\phi_{\lambda\alpha}^*(t)]dt\Big[\alpha s + (1-\alpha)\Pr\{s_j < s \,|\, a_j = \rho, \phi_{\lambda\alpha}^*\}\Big] \\ &= \alpha s \int_0^1 \phi_{\lambda\alpha}^*(t)dt + (1-\alpha)\int_0^s \phi_{\lambda\alpha}^*(t)dt \\ &\quad + \alpha s \int_0^1 [1-\phi_{\lambda\alpha}^*(t)]dt + (1-\alpha)\int_0^s [1-\phi_{\lambda\alpha}^*(t)]dt \end{aligned}$$

And the expected payoff from choosing $\rho$ is:

$$V_\rho^\alpha(s) = \alpha s \int_0^1 \phi_{\lambda\alpha}^*(t)dt + (1-\alpha)\int_0^s \phi_{\lambda\alpha}^*(t)dt + M\int_0^1 [1-\phi_{\lambda\alpha}^*(t)]dt$$

Using the logit specification for the quantal response function, we then apply logit choice probabilities to the difference in the expected payoff from $\phi$ and

[16] In a fully cursed equilibrium ($\alpha = 1$), all players choose strategies as if there is no correlation between the opponent's action and information. Thus, they all behave like a level 1 player in CH with random actions: $s_1^*(M) = M$.

[17] Note that player heterogeneity with respect to $\alpha$ will imply heterogeneity of cutpoints but it will still not solve the zero-likelihood problem: for any cursedness $\alpha \in [0, 1]$, it is always true that $s_\alpha^*(M) \leq M$. However, in our data set, we have many observations where players with strength $s > M$ choose $\rho$.

$\rho$ for each $s_i = s$. By inspection of $V_\phi^\alpha(s)$ and $V_\rho^\alpha(s)$, this difference is:

$$\Delta(s;\phi^*_{\lambda\alpha}) = \alpha s\int_0^1[1-\phi^*_{\lambda\alpha}(t)]dt+(1-\alpha)\int_0^s[1-\phi^*_{\lambda\alpha}(t)]dt-M\int_0^1[1-\phi^*_{\lambda\alpha}(t)]dt$$

and the $\alpha$-QRE in the simultaneous game is then characterized by:

$$\phi^*_{\lambda\alpha}(s) = \frac{e^{\lambda\Delta(s;\phi^*_{\lambda\alpha})}}{1+e^{\lambda\Delta(s;\phi^*_{\lambda\alpha})}} \quad \text{for all } s \in [0,1]$$

which can be solved numerically, for any value of $\alpha$.

In the sequential version of the game, we need to simultaneously solve for the first and second movers, $\phi^*_{\lambda\alpha 1}$ and $\phi^*_{\lambda\alpha 2}$, respectively. The expected payoff equations under $\phi$ and $\rho$ for the first mover are the same as in the simultaneous move game, so we have:

$$\begin{aligned}
V^\alpha_{\phi 1}(s_1) &= \alpha s_1\int_0^1 \phi^*_{\lambda\alpha 2}(s_2)ds_2 + (1-\alpha)\int_0^{s_1}\phi^*_{\lambda\alpha 2}(s_2)ds_2 \\
&\quad +\alpha s_1\int_0^1[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2 + (1-\alpha)\int_0^{s_1}[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2 \\
V^\alpha_{\rho 1}(s_1) &= \alpha s_1\int_0^1 \phi^*_{\lambda\alpha 2}(s_2)ds_2 + (1-\alpha)\int_0^{s_1}\phi^*_{\lambda\alpha 2}(s_2)ds_2 + M\int_0^1[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2
\end{aligned}$$

However, the expressions for the second mover are different, because expected payoffs are conditional on the observation that the first mover chose $\rho$:

$$\begin{aligned}
V^\alpha_{\phi 2}(s_2) &= \alpha s_2 + (1-\alpha)\frac{\int_0^{s_2}[1-\phi^*_{\lambda\alpha 1}(s_1)]ds_1}{\int_0^1[1-\phi^*_{\lambda\alpha 1}(s_1)]ds_1} \\
V^\alpha_{\rho 2}(s_2) &= M
\end{aligned}$$

So, the payoff differences for the first and second movers are, respectively:

$$\begin{aligned}
\Delta_1(s_1;\phi^*_{\lambda\alpha 2}) &= \int_0^1[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2\left[\alpha s_1+(1-\alpha)\frac{\int_0^{s_1}[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2}{\int_0^1[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2} - M\right] \\
\Delta_2(s_2;\phi^*_{\lambda\alpha 1}) &= \alpha s_2 + (1-\alpha)\frac{\int_0^{s_2}[1-\phi^*_{\lambda\alpha 1}(s_1)]ds_1}{\int_0^1[1-\phi^*_{\lambda\alpha 1}(s_1)]ds_1} - M
\end{aligned}$$

Note that the RHS of $\Delta_2$ is similar to the RHS of $\Delta_1$, except for the factor of $\int_0^1[1-\phi^*_{\lambda\alpha 2}(s_2)]ds_2$. Since this factor is smaller than 1, it means that the payoff differences to player 2 are magnified relative to player 1, which, in equilibrium, will result in $\phi^*_{\lambda\alpha 2}$ having higher slope and lower mean compared to $\phi^*_{\lambda\alpha 1}$. The two logit equilibrium conditions are:

$$\begin{aligned}\phi^*_{\lambda\alpha 1}(s_1) &= \frac{e^{\lambda\Delta_1(s_1;\phi^*_{\lambda\alpha 2})}}{1+e^{\lambda\Delta_1(s_1;\phi^*_{\lambda\alpha 2})}} \quad \text{for all } s_1 \in [0,1] \\ \phi^*_{\lambda\alpha 2}(s_2) &= \frac{e^{\lambda\Delta_2(s_2;\phi^*_{\lambda\alpha 1})}}{1+e^{\lambda\Delta_2(s_2;\phi^*_{\lambda\alpha 1})}} \quad \text{for all } s_2 \in [0,1]\end{aligned}$$

One can fit the logit version of the $\alpha$-QRE model to the dataset by finding the values of $\lambda$ and $\alpha$ that maximize likelihood of the observed frequencies of strategies. We estimated the best-fitting values by maximum likelihood for all four treatments and display the results in Table 5.

## 5.5 Model estimates

In this section we compare how well the different models fit the data, and explore the stability of the estimated parameters across the different treatments. We estimate the free parameters of all the models described in the previous section, for the treatments separately, pooling across the $M$-treatments and pooling across all treatments. For the QRE, CH and TQRE models, we consider both the behavioral strategy version and the cutpoint version. In all three models, the cutpoint version fit the data better than the behavioral strategy version in every single treatment and in all the pooled estimations. This is not surprising, given our earlier finding that most subjects exhibit choice behavior that is consistent with a cutpoint strategy. We therefore report in Table 5 only the results for the cutpoint versions of these models. We also refer only to these models in the discussion below.

| | | QRE-cut | | CH-cut | | TQRE-cut | | | $\alpha$-QRE | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | N | $\lambda$ | -lnL | $\tau$ | -lnL | $\gamma$ | $\tau$ | -lnL | $\lambda$ | $\alpha$ | -lnL |
| Sim .39 | 560 | 20.8 | 171.0 | 0.6 | 183.2 | 4.4 | 5.0 | 170.5 | 26.6 | 0.92 | 145.5 |
| Sim .50 | 560 | 11.3 | 213.6 | 0.3 | 211.7 | 449.0 | 0.4 | 210.3 | 18.4 | 0.77 | 202.9 |
| Sim All | 1120 | 16.2 | 387.8 | 0.5 | 397.4 | 6.9 | 2.7 | 386.3 | 21.3 | 0.85 | 355.6 |
| Seq .39 | 395 | 11.5 | 125.0 | 0.4 | 137.1 | 6.0 | 2.4 | 124.6 | 23.5 | 0.97 | 102.0 |
| Seq .50 | 386 | 9.3 | 140.5 | 0.5 | 137.9 | 142.0 | 0.5 | 136.6 | 15.8 | 0.75 | 138.2 |
| Seq All | 781 | 10.4 | 265.3 | 0.4 | 275.1 | 8.0 | 1.8 | 263.5 | 18.4 | 0.86 | 248.9 |
| All | 1901 | 13.0 | 656.8 | 0.5 | 672.6 | 10.0 | 1.8 | 651.2 | 20.1 | 0.85 | 605.9 |

**Table 5**. Model estimates.

There is substantial variation in the quality of fit across the different models. The better fitting models all converge to 0% for low strengths and to 100% for high strengths. The $\alpha$-QRE model, which generally fits the best of all these models, does not have cutpoints built into it explicitly, but boils down to a "soft" cutpoint model; that is, $\alpha$-QRE choice probabilities follow a logit function that equals .5 at an implicit cutpoint. Figure 4 displays the empirical and fitted $\phi$-probabilities as a function of strength. The empirical frequencies are aggregated into bins of 5 units of strength (1-5, 6-10, etc.) along the horizontal axis, with $\phi$-probabilities in the vertical axis. Fitted choice frequencies are based on out-of sample[18] parameter estimates for the cutpoint models and the $\alpha$-QRE model. All these models capture the upward sloping empirical frequency of $\phi$. All exhibit low and flat $\phi$ rates for low strengths and high and flat $\phi$ rates for high strengths.

[ Insert Figure 4 ]

The CH-cut and QRE-cut models fit the data equally well, in spite of major differences in the predicted fight curves. In the CH-cut model, level 0 players randomize over cutpoints, which implicitly endows level 0 players with a degree of rationality in the form of increasing fight probabilities. This model predicts quite well the very low fighting rates for strengths below 20%. QRE-cut does a poorer job for strengths in that regions and yet it generally fits better than CH-cut. This can be attributed to the fact that QRE-cut predicts that second movers will have sharper response functions than first movers, which is a feature of the data not captured by CH. TQRE-cut does not provide a substantial improvement over QRE-cut or CH-cut. In fact, the fitted $\phi$-rate for TQRE-cut and QRE-cut are very similar. They both share the problem of overestimating the fighting rates for subjects with low strength. The $\alpha$-QRE is the best fitting of all models, as it combines the elements of cursedness and stochastic choice. The pure cursed equilibrium predicts the steepest response of fighting probability as a function of strength. In fact, all players follow the same cutpoint strategy, which is a function of $\alpha$, the players' degree of cursedness.[19] Adding quantal response, produces a nice logit function of the fighting probability, that cuts .50 at $s \approx .40$ and which

[18]The displayed curves for the sequential treatments are constructed using the parameter estimates obtained from the pooled simultaneous data, and vice versa.

[19]This contrasts with the CH model where players with different levels of thinking follow different cutpoint strategies.

is consistent with the data. Furthermore, quantal response also introduces a steeper $\phi$ curve for the second movers than for the first movers, which is again consistent with the data.

There are some differences in fit between the .39 and the .50 treatments, with most models fitting the data from the .39 treatment better. It is hard to say exactly why this is the case, except to note that the empirical $\phi$ curves are steeper in the .39 data. There is virtually no difference in either the fit or the actual parameter estimates for the sequential and simultaneous treatments. The $\alpha$-QRE pooled estimates of $\lambda$ and $\alpha$ are not significantly different across the treatments, even at the 5% level, and the fit is identical (log Likelihood/N=-.318 in both cases).

### 5.6 Summary of estimation results

The main findings about the specific models we estimated are summarized as follows. First, all models capture the most basic qualitative properties of behavior: fight rates are increasing in $s$ and decreasing in $M$. Second, estimates are similar across $M$ treatments, and little power is lost by pooling treatments. Third, models based on QRE capture the fact that first movers behave differently from second movers. In particular, the $\phi$ function is steeper for second movers. Models based on CH capture the low fighting rates of players with strength below 20%. Fourth, the cutpoint versions of CH and QRE describe behavior better than the behavior strategy versions. This is not surprising, given the findings at the individual level that indicate widespread use of cutpoint strategies by our subjects. Fifth, TQRE-cut provides an almost identical fit than QRE-cut, suggesting that, in this game, the addition of hierarchical thinking to quantal response does not have a substantial impact. Last, the $\alpha$-QRE model fits best. The estimates of $\alpha$ are virtually identical for both the sequential and simultaneous games.

## 6 Conclusions

The compromise game is obviously very challenging to the cognitive abilities of players. This is true not only for our subjects but even for experienced microeconomists. In our experiment, players seem to understand some basic elements of the game, such as the cutpoint nature of the optimal strategy. However, they have problems figuring out the full logic of the unravelling argument. It is also important to realize that the game has some unusual properties: each player's action is relevant only if the rival does not play optimally, and the Nash equilibrium strategy is a best response only to the rival playing also the Nash equilibrium.

We conclude with some final comments. First, this paper only explores a few possible explanations for the surprising behavior we observed in these games of incomplete information. There are several other alternative approaches that might be interesting to explore as future research.

One candidate would be the social preference approach, which hypothesizes that models based on selfish preferences are mis-specified. However, it is not clear exactly what such models of pro-social behavior (fairness, reciprocity, or altruism) would imply in the present context. For instance, if we adopt the fairness formulation of Fehr and Schmidt (1999), the payoffs of winning, compromising and losing become $1 - \beta$, $M$, and $-\alpha$ respectively. Thus, if fairness considerations are sufficiently strong ($\beta \geq 1 - M$) agents always play $\rho$. Otherwise, agents want to set a lower cutpoint than their rival and the equilibrium again unravels to playing always $\phi$. Our subjects do not exhibit such extreme behavior. Reciprocity is also unlikely to account for the observed choices. The second movers fight more than the first movers. This means that observing $\rho$ increases rather decreases the willingness to reciprocate by responding also with $\rho$. Overall, the major problem with this game seems to be the high level of cognitive ability required to play optimally. Of course, the detailed implications of pro-social models are more complicated, once they are combined with stochastic choice behavior, which would be necessary for a careful evaluation of their statistical fit to the data.

A second direction to extend the models (including social preference models) is to explicitly allow for heterogeneity in the data. While the CH model is suggestive of heterogeneity, the attempts here an elsewhere to fit the model assumes homogeneity, since repeated observations of the same individual are treated as independent draws from the type space. In principle, one could extend the CH models to allow for fixed types. Similarly, the QRE and $\alpha$-QRE models could be extended to allow for heterogeneity with respect to $\lambda$ and $\alpha$, also with fixed types.

One of the most interesting findings is that the order of moves affects choices in this game. While we found one possible explanation for this phenomenon (QRE), there are probably other explanations that could be formalized, and we believe understanding this phenomenon more deeply constitutes an interesting challenge for future research. Indeed, a player's action is relevant only if the rival chooses $\rho$. Thus, first, second and simultaneous movers should all condition their strategy on that event. By contrast, the data shows that players who observe $\rho$ being played by their rival (second movers) respond more aggressively than players who must condition on the anticipation of that event (first movers). Even among subjects who do not observe the choice of the rival before playing, there is a difference between knowing that one's choices will be publicly observed before the rival makes his choice (first movers) and knowing that one's choices will not be observed (simultaneous movers). *Hypothetical conditioning on events seems to produce different behavior than observational conditioning on events.* Naturally, this has implications for many other games of asymmetric information, including common value auctions and voting behavior, where optimal choice requires voters to condition on being pivotal.

Last, we tried to give the equilibrium model a good shot at succeeding, by including a treatment where $M < .50$. In that treatment, fighting is ex ante fair and ex ante efficient. While subjects did respond to lower compromise payoffs by choosing higher $\phi$-rates, the increase was very slight. We conjecture that one would have to set very low $M$s in order to observe fighting rates close to the theoretical predictions. We also conjecture that $\rho$-rates would rocket if compromising was the socially efficient outcome ($M > .50$). This adds an interesting dimension to the problem since, as discussed in section 4.3, when $M > .67$ both players benefit when one of them is boundedly rational and the other is aware of the cognitive limitations of his rival.

# References


[1] Akerlof, G. (1970) "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics,* 84, 488-500.

[2] Camerer, C., T. Ho, and J. Chong (2004) "A Cognitive Hierarchy Model of Behavior in Games," *Quarterly Journal of Economics,* 119(3), 861-98.

[3] Camerer, C., T. Palfrey, and B. Rogers (2006) "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," *Working Paper,* California Institute of Technology, http://www.hss.caltech.edu/˜trp/hqre051206.pdf.

[4] Casella, A., A. Gelman, and T. Palfrey (2006) "An Experimental Study of Storable Votes," *Games and Economic Behavior,* in Press, Working Paper version available at http://www.hss.caltech.edu/˜trp/storable.final.pdf.

[5] Costa-Gomes, M. and V. Crawford (2006) "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, In Press.

[6] Crawford, V. and N. Iriberri (2005) "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" Working Paper, UCSD.

[7] Eyster, E. and M. Rabin (2005) "Cursed Equilibrium," *Econometrica*, 73: 1623-72.

[8] Fehr, E. and K. Schmidt (1999) "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, 114(3), 817-868.

[9] Forsythe, R, R.M. Isaac, and T. Palfrey (1989) "Theories and Tests of 'Blind Bidding' in Sealed Bid Auctions," *Rand Journal of Economics*, 20, 214-238.

[10] Goeree, J., C. Holt, and T. Palfrey (2005), "Regular Quantal Response Equilibrium," *Experimental Economics*, 8, 347-367.

[11] Kagel, J and D. Levin (2002) *Common Value Auctions and the Winner's Curse*, Princeton University Press: Princeton, NJ.

[12] Lynch, M., R. Miller, C. Plott, and R. Porter (1984), "Product Quality, Consumer Information, and 'Lemons' in Experimental Markets," in *Empirical Approaches to Consumer Protection Economics,* (P. Ippolito and D. Scheffman eds.), FTC: Washington D.C., pp. 251–306.

[13] McKelvey, R. and T. Palfrey (1995) "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 10, 6-38.

[14] McKelvey, R. and T. Palfrey (1998) "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1, 9-41.

[15] Milgrom, P. and N. Stokey (1982) "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17-27.

[16] Morris, S (1994) "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327-1347.

[17] Myerson, R.B. and M.A. Satterthwaite (1983) "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, 29(2), 265-281.

[18] Nagel, R. (1995) "Unraveling in Guessing Games: An Experimental Study," *American Economic Review,* 85(5), 1313-1326.

[19] Palfrey, T. and J. Prisbrey (1996) "Altruism, Reputation, and Noise in Linear Public Goods Experiments," *Journal of Public Economics*, 61, 409-27.

[20] Samuelson, W. and M. Bazerman (1985) "The Winner's Curse in Bilateral Negotiations," in *Research in Experimental Economics* (V. Smith ed.), JAI press, 105-137.

[21] Sonsino, D., I. Erev, and S. Gilat (2001) "On Rationality, Learning and Zero-Sum Betting – An Experimental Study of the No-Betting Conjecture," Working Paper, Technion.

[22] Sovic, Y. (2004) "Strength of Dominance and Depths of Reasoning," Working Paper, University of Oslo.

[23] Wang, S. (2006), "Misunderstanding the Other," *Term Paper*, Princeton University.

# Appendix: Sample Instruction Script

Thank you for agreeing to participate in this research experiment on group decision making. During the experiment we require your complete, undistracted attention. So we ask that you follow these instructions carefully. You may not open other applications on your computer, chat with other students, or engage in other distracting activities, such as using your cell phones or head phones, reading books, etc.

For your participation, you will be paid in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. So it is important that you listen carefully, and fully understand the instructions before we begin. You will be asked some review questions after the instructions, which have to be answered correctly before we can begin the paid session.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiment except according to the rules described in the instructions.

We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you privately.

This experiment will begin with a brief practice session to help familiarize you with the rules. The practice session will be followed by a paid session. You will not be paid for the practice session.

This paid session of the experiment has 2 parts. In each part you will make choices over a sequence of 20 different decision rounds so in total you will make 40 decisions. In each round, you will receive a payoff, that depends on your decision that round and on the decision of one randomly selected participant you are matched with. We will explain exactly how these payoffs are computed in a minute.

At the end of the paid session, you will be paid the sum of what you have earned in all 40 decision rounds, plus the show-up fee of $10.00. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in POINTS. Your DOLLAR earnings are determined by multiplying your earnings in POINTS by a conversion rate. In this experiment, the conversion rate is 0.006, meaning that 100 POINTS equals 60 cents.

Here is how each decision round, or match, works. First, the computer randomly matches you into pairs. Since there are 14? (note to reader: depends on session) participants in today's session, there will be 7? (note to reader: depends on session) matched pairs in each decision round. You are not told the identity of the participant you are matched with. Your payoff depends only on your decision and the decision of the one participant you are matched with. What happens in the other pairs has no effect on your payoff and vice versa. Your decisions are not revealed to participants in the other pairs.

Next, the computer randomly assigns a number to you, which is equally likely to be any number between 1 and 100. This number is called your "strength." Each strength number is chosen independently for each participant. Therefore usually you and the person you are matched with will have different numbers, although there is a very small (1%) chance the other participant in your pair has the same strength you have. You are told your strength, but will not be told the strength of the other participant until after you have made your decision.

You then have to make a decision to take one of two possible actions. These two actions are called "fight" and "retreat". If both of you choose retreat, then both of you will receive a payoff of 40 points each. However, if either of you chooses fight, then the one with the greater strength receives a 95 points payoff and the one with less strength receives a 5 points payoff. Ties are broken randomly.

[SCREEN 1] This slide shows a summary of the Payoffs

Each of you must make your decision to fight or retreat at the same time, so neither of you are told what the other participant chose (or their strength) until after both of you have made your choices. The match is over when you and the person you are matched with have both made a decision, and the computer will show you the results of your match only.

When all pairs have finished the match and seen the results, we proceed to the next match. For the next match, the computer randomly reassigns all participants to a new pair, and randomly reassigns a new strength to each participant. Your new strength assignment does not depend in any way on the past decisions or strengths of any participant including yourself. Strength assignments are completely independent across pairs, across participants, and across matches. After learning your new strength assignment, you choose either "fight" or "retreat" and receive payoffs in a similar manner as in the previous match.

This continues for 20 matches, at which point Part 1 of the experiment is over. I will read you the instructions for Part 2 after we complete Part 1.

We will now begin the Practice session and go through two practice rounds. During the practice matches, please do not hit any keys until you are asked to, and when you enter information, please do exactly as asked. Remember, you are not paid for these 2 practice rounds. At the end of the second practice round you will have to answer some review questions. Everyone must answer all the questions correctly before the experiment can begin.

[AUTHENTICATE CLIENTS]

Please double click on the icon on your desktop that says "c". When the computer prompts you for your name, type your First and Last name. Then click SUBMIT and wait for further instructions.

[START GAME]

[SCREEN 2]

You now see the first screen of the experiment on your computer. It should look similar to this screen.

At the top left of the screen, you see your subject ID. Please record that on your record sheet now. You have been randomly matched by the computer with exactly one of the other participants. This pair assignment will change after each match.

You have been assigned your strength for this match, which is revealed to you on your screen. [point on overhead]. Your exact strength number on your own screen would probably be different from the one on this slide.

The participant you are matched with was also randomly assigned a strength, but that will not be revealed to you until the end of the match. All you know now is that their strength is some number between 1 and 100, with every number being equally likely.

There are two buttons, one marked “Fight” and one marked “Retreat”. You must choose one of those two buttons, but please do not do so yet. I want to remind you how your payoffs will be computed. If you and the person you are matched with BOTH choose retreat, then each of you receives a 40 points payoff. If either one of you chooses fight, then whoever has the higher strength receives 95 points and whoever has the lower strength receives 5 points. If you have the same strength, then the computer will randomly choose one of you to receive 95 points and the other to receive 5 points.

At this time, if your subject ID is even, please click on the button labelled “fight”. If your subject ID is odd, please click on the button labelled “retreat”.

When everyone has made a choice, you are told the choice made by the participant you are matched with and also told that participant’s strength. The outcome is summarized on your screen. [show on overhead screen]

[SCREEN 3]

Each Round Summary is shown on the center of the screen.

The bottom half of your screen contains a table summarizing the results for all matches you have participated in. This is called your history screen. It will be filled out as the experiment proceeds. Notice that it only shows the results from your pair, not the results from any of the other pairs. PLEASE record this information on your record sheet.

We now proceed to the next match.

For the next match you will be randomly re-matched into pairs, and randomly receive new strength assignments.

[START next MATCH]

Please notice your new strength assignment. [Reader: Ask if everyone sees it, and wait for confirmation from them.] Please make the opposite decision in match 2 than you made in match 1. That is, if your subject ID is even please click on the retreat button and if your subject ID is odd, please click on the “fight” button, and then wait for further instructions.

[wait for them to complete match 2]

Practice match 2 is now over.

Please complete the review questions before we begin the paid session. Once you answer all the questions correctly, click submit. After both participants in your pair have answered the first round of questions, the next round of questions will appear. Please answer all questions correctly and click submit and the quiz will disappear from your screen. [WAIT for everyone to finish the Quiz]

Are there any questions before we begin with the paid session? We will now begin with the 20 paid matches of Part 1. Please pull out your dividers for the paid session of the experiment. If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you.

[START MATCH 2]

[After MATCH 21 read:]

We have now reached the end of Part 1. Your total payoff from this part is displayed on your screen. Please record this on your record sheet and CLICK OK. We will now give you instructions for Part 2. Please listen carefully.

Part 2

Part 2 of the experiment will take place over a sequence of 2 practice and 20 paid matches. This Part is almost exactly the same as Part 1, with one difference. For each pair, one of the participants will choose to retreat or fight before the other participant makes a choice. The other participant will then be told the first participant's decision and will then make their decision in response. Payoffs are computed exactly as before. For each match, the computer randomly selects the participant to decide first, so sometimes you will decide first and sometimes you will decide second. The assignment of who decides first or second does not in any way depend on the strength assignment or past decisions. The computer program just randomly assigns one participant for each pair to decide first.

We will now proceed through two practice matches to familiarize you with the screens, which are slightly different than Part 1.

[Go to Match 22]

You now see the first screen of the experiment on your computer.

In this part of the experiment, the computer randomly assigns an order in which the two members of your pair make decisions. If you are assigned as the first decision maker in your match, your screen should look similar to this:

[SCREEN 4]

[Describe the screen by pointing and READ THE SCREEN]

Of course your exact strength number on your own screen would probably be different from the one on this overhead.

If you are assigned as the second decision maker in your match, your screen looks like this:

[SCREEN 5]

[Describe the screen by pointing and READ THE SCREEN]

Of course your exact strength number on your own screen would probably be different from the one on this overhead.

The first decision maker in each pair must make a decision. If you are decision maker one and your ID is even please click on the Fight button now. If you are decision maker one and your ID is odd please click on the Retreat button now. If you are decision maker two, please wait until it is your turn to make a decision.

Next, the first decision maker's choice is revealed to the second decision maker. Please do not make any decisions until I finish explaining. The screen looks like:

[SCREEN 6] if decision maker one chose fight, READ the SCREEN

[SCREEN 7] if decision maker one chose retreat READ the SCREEN

After viewing this information, the second decision maker is prompted to a choice. If decision maker one chose Fight, then the outcome does not depend on decision maker two's choice. In this case we simply ask decision maker two to click on the "continue" button. If decision maker one chose retreat, then the outcome does depend on decision maker two's choice, so decision maker two must now make a choice of fight or retreat.

This information is summarized on this slide

[SCREEN 8]

If you are decision maker two and you have a choice, if your ID is even, please click the Retreat button now. If your ID is odd, please click on Fight now. Otherwise, please click on the continue button now. This is important so please do not forget to do so. The match cannot proceed until the second decision maker has clicked a button.

The results of the match are then displayed for both decision makers in the pair. The screen should look like:

[SCREEN 9] for first decision maker

[SCREEN 10] and like this for second decision maker

We will now proceed to the second practice match [CLICK NEXT MATCH]. When you are prompted to make a decision, please make the opposite decision from your decision in the first practice match. That is, if your ID is even, click on Retreat, and if your ID is odd, click on Fight. Please go ahead and make your choices. If you are decision maker two and you see the continue button, please remember to click on it or it will delay the experiment.

[Advance to match 23 and wait for participants to finish]

The practice match is now over. Are there any questions before we begin the 20 paid matches?

[SCREEN 11] Here is a summary of the Payoff

[START MATCH 24]

[After MATCH 44, read:]

Your Total Payoff for both parts is displayed on your screen. Please record this payoff on your record sheet and remember to CLICK OK after you are done.

[CLICK ON WRITE OUTPUT]

Your total payoff is this amount plus the show-up fee of $10. We will pay each of you in private in the next room in the order of your Subject ID number. Remember you are under no obligation to reveal your earnings to the other players.

Please put the mouse behind the computer and do not use either the mouse or the keyboard at all. Please remain seated and keep the dividers pulled out until we call you to be paid. Do not converse with the other participants or use your cell phone while in the laboratory.

Thank you for your cooperation.

Could the person with ID number 0 go to the next room to be paid.

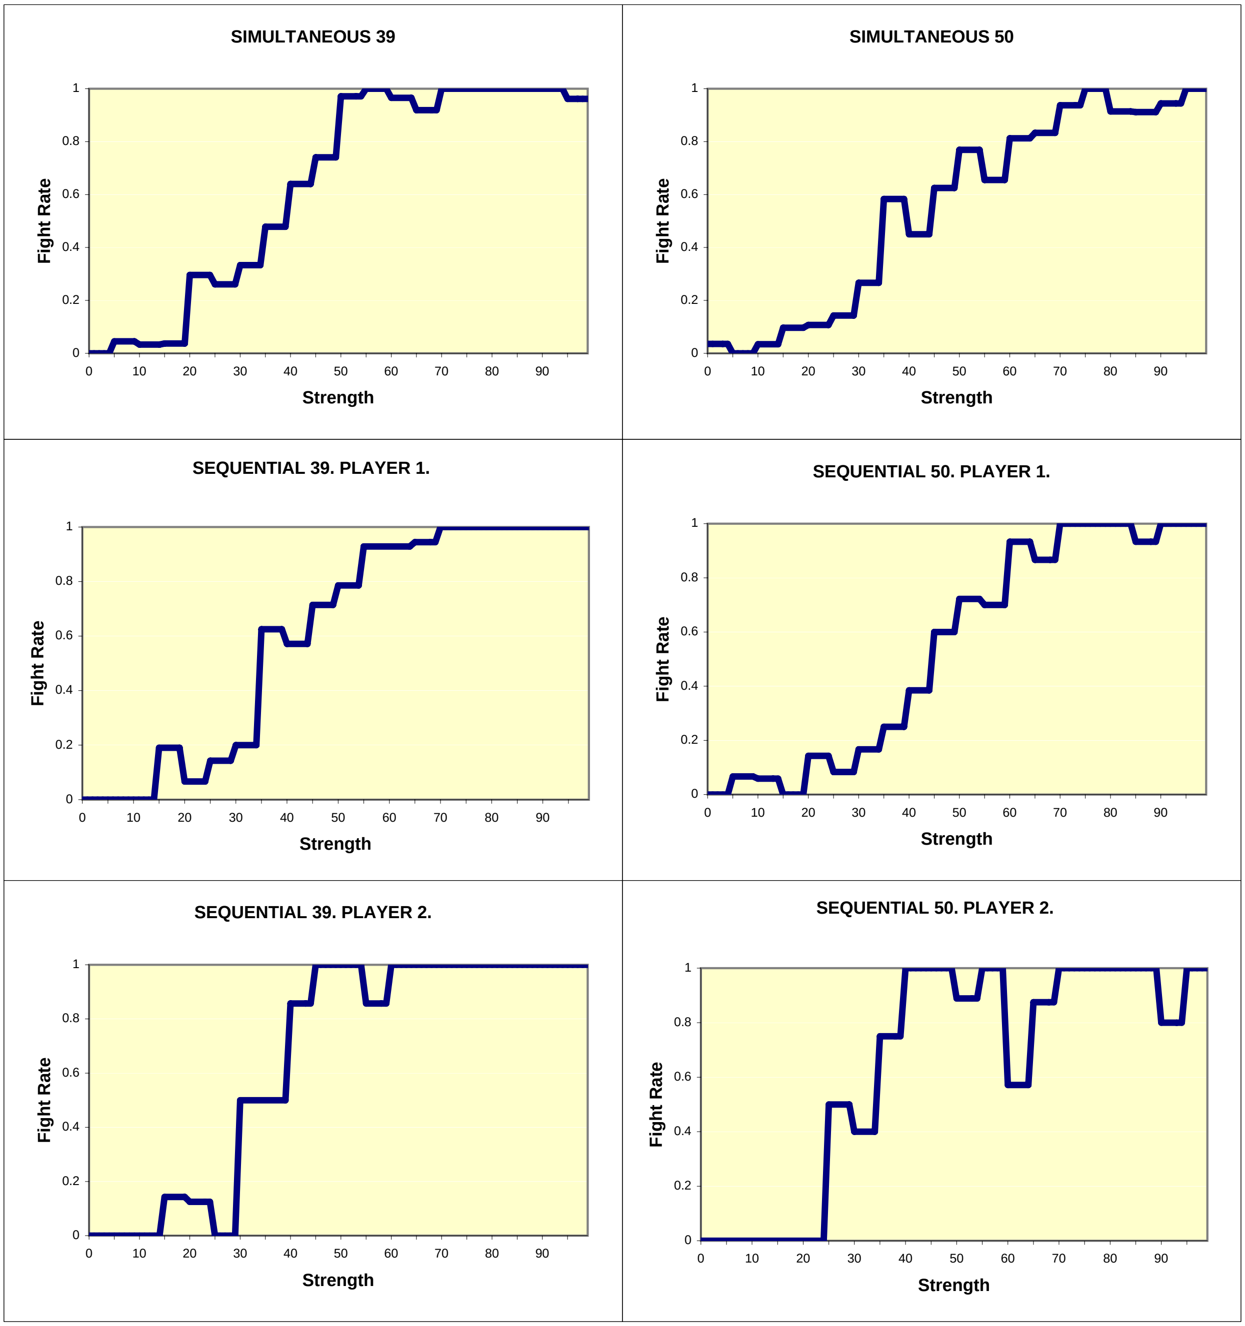


Figure 1. Empirical Fight Rates.

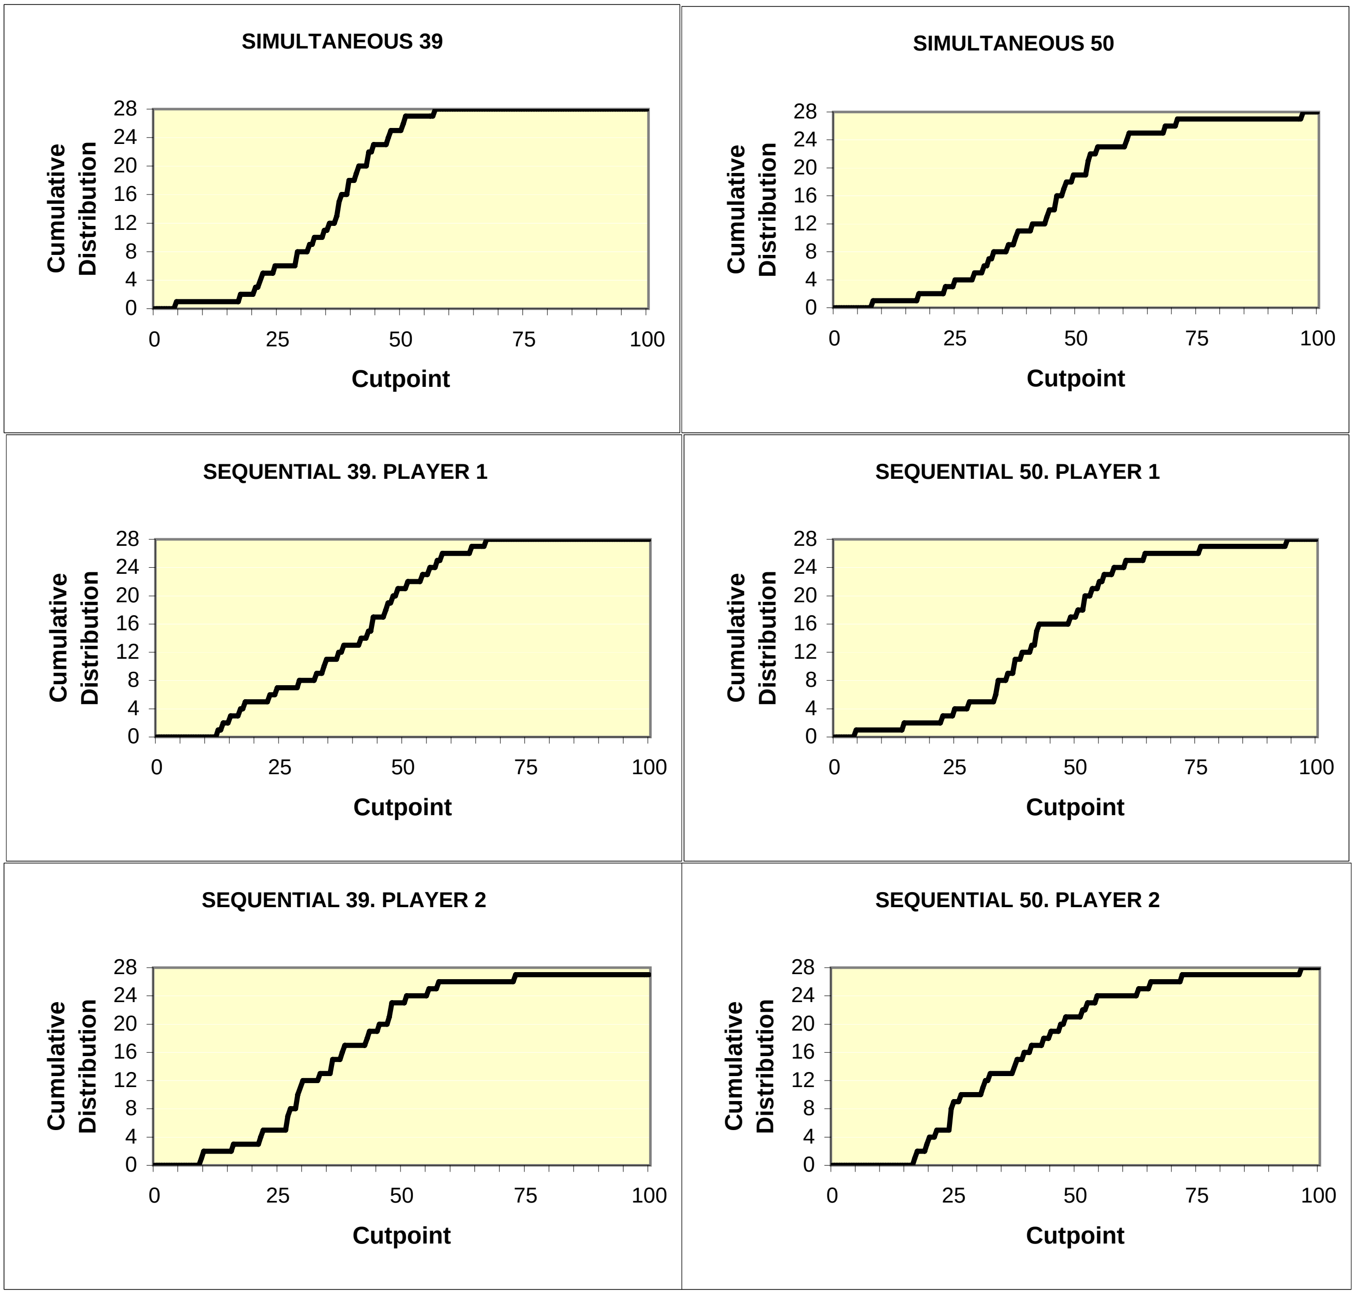


Figure 2.  Distribution of cutpoints, by condition.

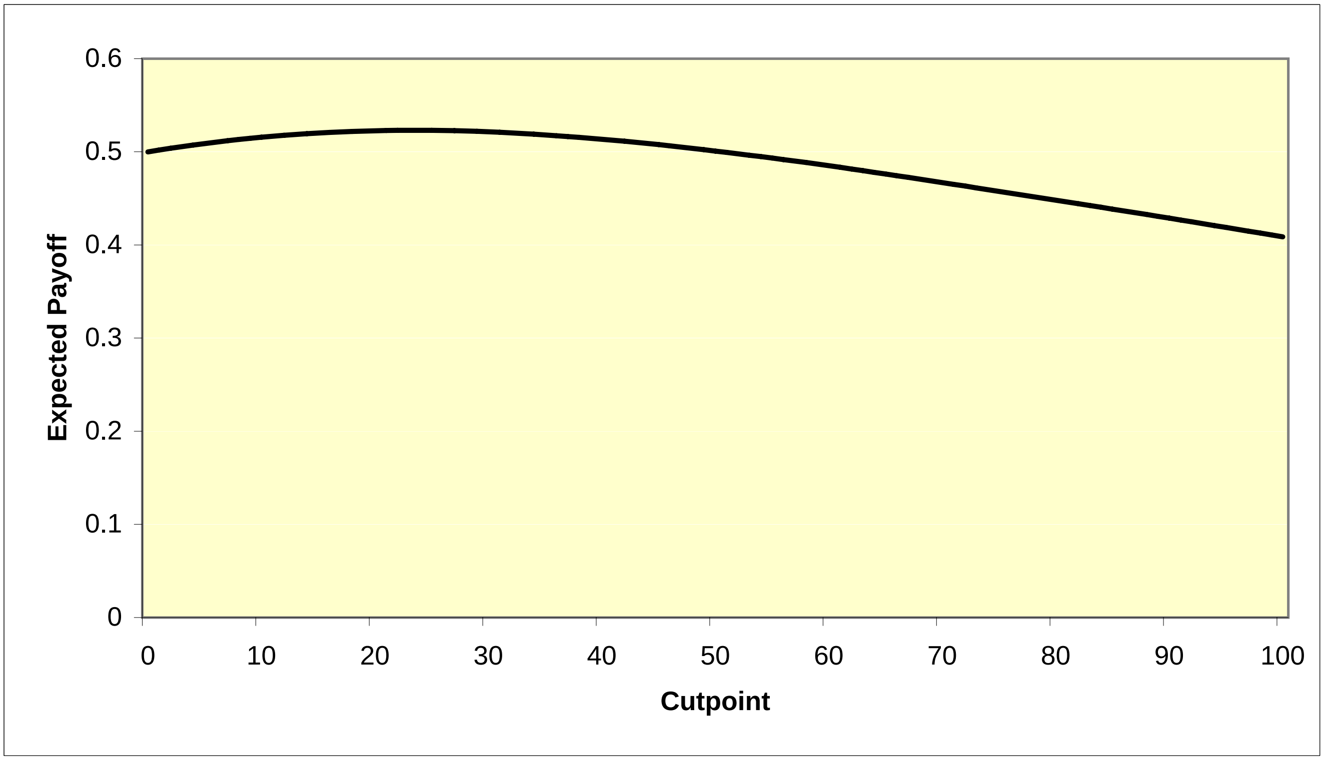


Figure 3. Expected payoff for different cutpoints against empirical fight rates. (Simultaneous 50)

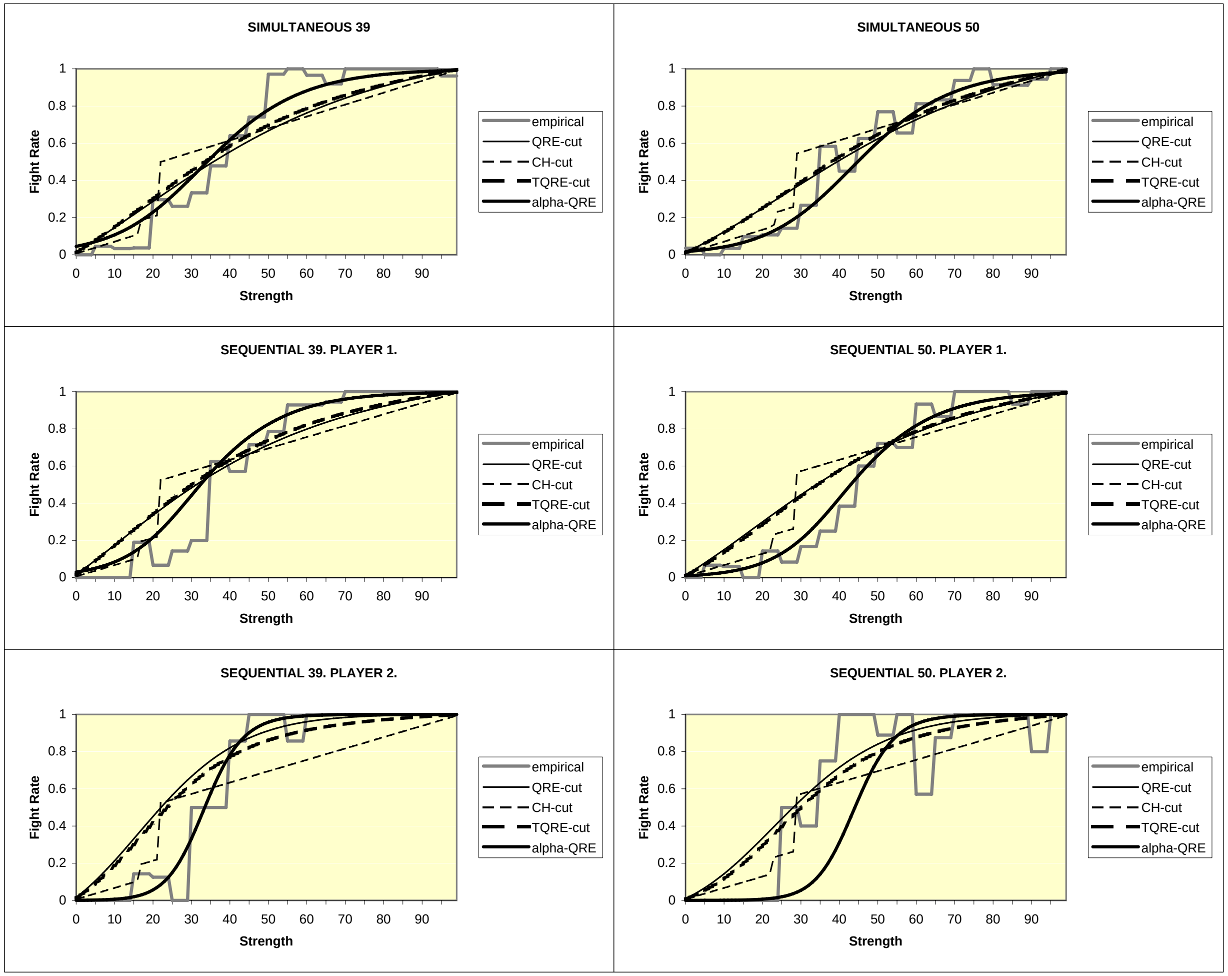


Figure 4. Empirical and Fitted Fight Rates.