

**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

April 9, 2008

*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, Universidad Carlos III, Universitat Autònoma de Barcelona, University of Pennsylvania, the Fundación Ramon Areces Conference on Experimental and Behavioral Economics in December 2006, and 2006 ESA meeting in Tucson for comments, and Shivani Nayyar, Stephanie Wang, and Rumen Zarev for research assistance. Corresponding addresses: Juan D. Carrillo, Department of Economics, University of Southern California, Los Angeles, CA 90089-0253, <juandc@usc.edu>. Thomas R. Palfrey, Division of the 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 compromise. 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 compromise, conflict is avoided and each receives an intermediate payoff. The only equilibrium in both the sequential and simultaneous versions of the game is for players to always fight, independent of their own strength. In our experiment, we have two main treatments: whether the game is played simultaneously or sequentially; and the magnitude of the intermediate payoff. We observe: (i) frequent compromise; (ii) significantly more fighting the lower the compromise payoff; (iii) significantly less fighting by first than second movers; and (iv) almost no evidence of learning. We explore several theories of cognitive limitations in an attempt to understand the reasons underlying these anomalous findings, including quantal response equilibrium, cognitive hierarchy, and cursed equilibrium.

JEL classification: C92, D82.

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

1 Introduction

A major insight 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 the other side of the coin, namely a situation where exchanges (or other type of agreements) are not mutually beneficial and ask the following question: can private information induce agents to reach agreements that one of them will (ex post) regret? An 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. 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, it is obvious that the strongest nation has always an incentive to attack the weakest one. Thus, a war is inevitable. More interestingly, the equilibrium is also war if the leader of each nation knows its own military strength but knows only the probability distribution of the other nation's strength (and therefore is uncertain over his chances of winning). This would be true, for example, even if the benefits of winning the war were only slightly greater than the peace benefits, the costs of losing the war were enormous, and the uncertainty about the other nation's strength were large. The logic is much like the unraveling 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

reach an agreement that cannot ex post benefit both parties: litigation, electoral debates, and firm competition.

We report here an experiment that analyzes behavior in several variations of this two-sided asymmetric information environment in the laboratory. In all the variations, the equilibrium outcome predicted by the theory is the same: fighting ensues with probability one. We obtain three main results, which are inconsistent with standard game theory. First and foremost, fighting occurs much less often than predicted by theory. Rather than 100%, we observe fight rates in the range of 50%-70%. The outcome characterized by both agents compromising arises with surprising frequency, nearly one-quarter of the time in some sessions. In terms of our example it means that, contrary to the predictions of game theory, a war can be avoided if the military strengths of countries are privately (rather than publicly) known. Second, fight rates are affected by the compromise payoff. In both the sequential and simultaneous treatments, agents are less likely to fight the higher the compromise payoff. Third, in the sequential version, the strategies of first and second movers are different in two ways: second movers are more likely to fight than first movers, and the behavior of second movers is more responsive to strength and less erratic than that of first movers.

We also obtain some findings about individual behavior. Individual choice is consistent with the use of cutpoint strategies: fight if and only if strength is above a certain critical threshold. However, instead of the cutpoints being at (or at least close to) the minimum strength, as predicted by the theory, we find that players use cutpoints in an intermediate range. The use of cutpoint strategies indicates that subjects have some understanding of the game, and the source of violations of equilibrium has to do with the cognitive difficulty to choose the cutpoint optimally. Perhaps more interesting, we find substantial heterogeneity in the choice of cutpoints across subjects, and also important differences in the distribution of cutpoints across treatments. Finally, all the results are robust with respect to experience, that is, there is little evidence of learning.

We then explore three recent theories of cognitive limitations in games, and analyze the data to investigate the extent to which the insights from these alternative theories can account for these anomalies. The three approaches we explore are *equilibrium stochastic choice*, *levels of strategic sophistication*, and *naïve belief formation*. 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). The estimated parameters for each model are relatively constant across treatments. However, there are some important differences in the predictions, that lead to differences in the fit of the models. An important result is that only the QRE model captures the tendency of second movers to fight more often than first movers. The CH and CE models capture the aggregate tendency of players to fight with probability close to one when their strength is sufficiently high and with probability close to zero when their strength is sufficiently low. The CH model predicts that the distribution of individual cutpoints will be multimodal, clustered around three or four numbers, which is not reflected in the data. Finally, the best fit is obtained with a hybrid of QRE and CE, which combines cursedness and stochastic choice.

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.

2.1 Some introductory examples

Consider two agents who must decide whether to split a surplus in a pre-specified manner (compromise) or try to reap all the benefits (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 situations fit this general description, in addition to the example of international conflict described in the introduction. 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 an electoral campaign, each candidate can drive its rival into a public debate, where some qualities of the contenders are revealed to voters, which affects the electoral outcome. In the absence of a debate, voters must rely on expected qualities. In a product market competition, firms offering horizontally differentiated products may start an R&D race. The winner monopolizes the market and the probability of winning

is proportional to the privately known quality of their research department. Alternatively, firms can avoid the race and share the market. In all these cases, there are only two possible outcomes: settlement, peace, no debate, market sharing vs. trial, war, debate, R&D race. The first outcome needs the agreement of both players whereas each player 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. Also, the utility of agents under the different outcomes may depend on some exogenously given private information parameters. 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. More generally, in the no-agreement outcome, there is always one (ex post) winner and one (ex post) loser relative to the agreement outcome. The total surplus from compromise varies across these different applications, although we show below that the equilibrium does not depend on the compromise payoffs. Wars are typically costly and socially inefficient, so there is as peace dividend. Litigation also involves a waste of resources when compared to early settlements (but not compared to last minute settlements). Electoral debates are roughly neutral if candidates are only interested in winning the election. However, they provide information regarding the merits of different proposals, which can be valuable if candidates care also about policy outcomes. As for market competition, the profits of a monopolist are typically higher than the sum of profits of two duopolists, which suggests that firms' surplus is increased when they fight for supremacy rather than splitting the market.¹

2.2 Formalizing the game

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 ρ stands for "retreat" and ϕ 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

¹In these examples, we are not including the welfare of third parties such as society, voters or consumers. These are also likely to be different across outcomes.

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, depending on the context, the socially efficient action may be compromise or no-compromise or it may even be a zero-sum game: $\alpha_i(s_1, s_2) + \gamma_j(s_1, s_2) \gtrless \beta_i(s_1, s_2) + \beta_j(s_1, s_2)$ for all $s_i > s_j$.

2.3 Equilibrium

Given this structure, we can analyze the Perfect Bayesian Equilibrium (PBE) for the sequential version of the game where agent 1 moves first and agent 2 moves second. We have the following result.

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

Proof. Suppose that there exist two sets $\tilde{S}_1 \subseteq S_1$ and $\tilde{S}_2(\tilde{S}_1) \subseteq S_2$ 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 $\bar{s}_1 = \max_{s_1 \in \tilde{S}_1}$ and $\bar{s}_2 = \max_{s_2 \in \tilde{S}_2(\tilde{S}_1)}$.

According to this PBE, once agent 2 has observed $a_1 = \rho$, the following inequality must be satisfied:

$$\begin{aligned} \int_{s_1 \in \tilde{S}_1} \beta_2(s_1, s_2) dF_1(s_1 | s_1 \in \tilde{S}_1, s_2) &\geq \int_{s_1 \in \tilde{S}_1 \cap s_1 < s_2} \alpha_2(s_1, s_2) dF_1(s_1 | s_1 \in \tilde{S}_1, s_2) \\ &+ \int_{s_1 \in \tilde{S}_1 \cap s_1 > s_2} \gamma_2(s_1, s_2) dF_1(s_1 | s_1 \in \tilde{S}_1, s_2) \quad \forall s_2 \in \tilde{S}_2(\tilde{S}_1) \end{aligned}$$

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 $s_2 = \bar{s}_2$. Since $\alpha_2(s_1, s_2) > \beta_2(s_1, s_2) > \gamma_2(s_1, s_2)$, the inequality necessarily implies that $s_1 < \bar{s}_2$ must be binding at least for some $s_1 \in \tilde{S}_1$. Therefore, $\bar{s}_2 < \bar{s}_1$. 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:

$$\begin{aligned} \int_{s_2 \in \tilde{S}_2(\tilde{S}_1)} \beta_1(s_1, s_2) dF_2(s_2 | s_1) &\geq \int_{s_2 \in \tilde{S}_2(\tilde{S}_1) \cap s_2 < s_1} \alpha_1(s_1, s_2) dF_2(s_2 | s_1) \\ &+ \int_{s_2 \in \tilde{S}_2(\tilde{S}_1) \cap s_2 > s_1} \gamma_1(s_1, s_2) dF_2(s_2 | s_1) \quad \forall s_1 \in \tilde{S}_1 \end{aligned}$$

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

Using the same reasoning as before, $\bar{s}_1 < \bar{s}_2$. 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. \square

The intuition is simple. In this class of games, agents know that good news for them is bad news 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 it, and vice versa. 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. Last, the result is also unchanged if agents play simultaneously. Indeed, 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, this does not make any difference since, both in the sequential and the simultaneous versions, each agent knows that his action is only relevant if the rival offers a compromise. Thus the outcome of the Bayesian Nash Equilibrium (BNE) is, just like for the PBE, always no-compromise. This result is summarized as follows.

Corollary 1 *The outcome of the game is still "no-compromise" if agents are risk-averse and if they announce their strategy simultaneously.*

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", ϕ , or "retreat", ρ . If 1 chooses ϕ , 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 ρ ,

then it is agent 2's turn. If agent 2 chooses ϕ , 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 ρ , 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 via 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.

In our design, the total surplus under compromise is either smaller than ($M = .39$) or equal to ($M = .50$) the total surplus under no compromise. As discussed earlier, this specific choice of parameters fits some examples (firm competition, electoral debate) better than others (military conflict, litigation). There were at least two reasons for focusing on these values. First, the $M = .50$ is a natural benchmark, corresponding to a constant sum game, where there is no difference in the efficiency of the fight and compromise outcomes. Second, we wanted to choose another value of M because the various behavioral theories we were testing all predict a negative comparative static effect of M on the probability of fighting. Changes in either direction would allow us to test this. Our choice of a lower value of M was made because we thought it was the more interesting direction, since for $M < .5$ always fighting is not only still the unique Nash equilibrium but is also ex ante efficient and ex ante fair. This allows us to better distinguish social preference or fairness-based explanations for excessive compromise (see section 5.5) from behavioral models based on cognitive limitations (see sections 5.1 to 5.4), and also gives Nash equilibrium its best shot. If the excessive compromise disappears, this would lend support for social preference theories based on ex ante fairness and efficiency, whereas, if the excessive compromise persists, it lends support for the models of cognitive limitations.

3.2 Relation to the experimental literature

Some related games have been studied in the laboratory. Below we describe (from most to least similar) two simultaneous games of multi-sided asymmet-

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

ric 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. As in our game, all the BNE imply no trade. There are three main differences. First, the risky outcome requires agreement in the betting game whereas the safe outcome requires agreement in the compromise game. Second, the information structure is simpler than in ours: there are only four possible states, and in two of them one agent has full information. The common knowledge of this information partition triggers very naturally unraveling to no-trade. This special partition is also likely to facilitate learning. Third, a sequential version of the betting game has not been studied.

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.

Adverse selection game (Akerlof, 1970). This game also predicts some unraveling. 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. Holt and Sherman (1994) prove that buyers may underbid or overbid depending on the treatment conditions. 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 determines whether an informed seller reveals the quality of his good to an uninformed buyer. 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 in our game, best response dynamics predicts unraveling for a wide range of parameters. Since the beauty contest

is a static game of complete information, the details of convergence are different. Even the most naïve learning rules (such as ‘play optimally given the outcome in the past round and assuming that nobody else revises his strategy’) predict rapid convergence if the beauty contest game is played repeatedly. The experimental data confirms this prediction, and a natural question is whether a similar convergence pattern is found in the compromise game.

3.3 Experimental design and procedures

We conducted five sessions with a total of 56 subjects, using a simple 2×2 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

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

⁵In 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.

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 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 ϕ 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,

⁶In one of the sessions, the sequential version was played in rounds 1 – 20 and the simultaneous version was played in rounds 21 – 40. We found no significant order effects.

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 fight vs. retreat, without conditioning on the actual draws of s_i . Table 2 shows the relative fight rates in the experiment, broken down by compromise payoff and order of moves, with the number of observations in parenthesis.

Order	Position	$M = .39$	$M = .50$	Pooled
Sequential	First	.589 (280)	.538 (264)	.564 (544)
Sequential	Second	.643 (115)	.566 (122)	.603 (237)
Sequential	Both	.605 (395)	.547 (386)	.576 (781)
Simultaneous	—	.657 (560)	.573 (560)	.616 (1120)
Pooled		.636 (955)	.562 (946)	.599 (1901)

Table 2. Unconditional fight rates.

There are several differences in fight rates across treatments and conditions.

First, there is a clear difference between ϕ rates in the $M = .39$ and the $M = .50$ treatments. Fighting is chosen approximately 13% more frequently when the compromise dividend is lower, and this difference is statistically significant ($p < .01$).⁷ This difference is observed in both the sequential and simultaneous treatments, for both the first and second movers separately, and these differences are all approximately the same magnitude. In looking at subsamples based on treatment or condition, the lower significance levels reflect the relatively small number of observations in each treatment: $p = .05$ level for the simultaneous treatment (most observations); $p = .10$ level for the sequential treatment; and $p > .10$ for first movers and second movers, separately.

⁷The significance levels reported here and later in the paper are based on standard tests treating the observations as independent. Later in the paper, we consider individual effects and learning.

Second, first movers in the sequential game fight less frequently than second movers, both in the .39 and the .50 treatments. The differences in means are not statistically significant.

Third, there is more fighting in the simultaneous treatment than in the sequential treatment. These results are mirrored in a simple probit dummy variable regression, where the dependent variable is whether a subject chose to fight and the independent variables are M , "sim" (= 1 for the simultaneous treatment), and "role" (= 1 for second movers). The estimated coefficients and standard errors are given in table 3.

	Coefficient	St. error	t -stat.
Constant	1.023	0.267	3.83
M	-0.019	0.006	-3.29
sim	0.134	0.066	2.03
role	0.106	0.099	1.08

Table 3. Probit regression of fight rates.

The logic of the game suggests that, over time, learning should lead to unraveling. 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 ϕ 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 4 presents the ϕ rates, broken down by experience level. The number of observations is in parenthesis.

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 4. Unconditional fight rates by experience level.

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

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 ϕ rates are not a sufficient statistic for the actual strategies. A behavior strategy in each game is a probability of choosing ϕ 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 frequencies are aggregated into bins of 5 units of strength (1-5, 6-10, etc.) along the horizontal axis, with ϕ -probabilities in the vertical axis.

[FIGURE 1 HERE]

These graphs suggest that second movers in the sequential version of the game behave differently from first movers in at least two ways. Second movers fight more than first movers. 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. These results are also supported by a probit regression similar to the one reported in table 3, but including the independent variable s to control for strength. The coefficients and t -statistics are

reported in table 5. Two important new results emerge when controlling for strength. First, the difference in fight rates between the simultaneous and sequential treatments, which was significant without controlling for s , is no longer significant, indicating that the earlier finding was due to the different realizations of s in the two samples. Second, the coefficient on "role" is now marginally significant, suggesting that the variation in sample draws of s obscured the differences between the fight rates of first and second movers.

	Coefficient	St. error	t -stat
Constant	-0.487	0.376	-1.296
M	-0.041	0.008	-4.906
sim	0.076	0.093	0.818
role	0.229	0.139	1.640
s	0.054	0.002	26.364

Table 5. Probit regression of fight rates, controlling for strength

The figure also suggests that second movers 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 curves for player 2 in figure 1. Table 6 describes aggregate behavior by treatment and position via a probit model where the independent variables are strength and a constant term. The steeper response by the second mover is confirmed for the .39 treatment, where the slope of the second mover's response curve is significantly steeper than the slope of the first mover. There is no significant difference in the slope coefficient for the .50 treatment.

Position	M	Constant	Slope	-lnL	% pred.	#obs.
Seq. 1	.39	-2.73 (.307)	.068 (.0071)	76.7	88.6	280
Seq. 2	.39	-3.11 (.585)	.088 (.0016)	21.6	92.2	115
Seq. 1	.50	-2.53 (.283)	.056 (.0057)	83.4	87.1	264
Seq. 2	.50	-1.83 (.162)	.044 (.0063)	45.6	86.9	122

Table 6. Probit comparing slope of response for first and second movers.

4.3 Individual cutpoint analysis

In order to address the question of conditional fight rates, treatment effects, and differences between first and second movers more carefully, we turn to an analysis of individual choice behavior. If subjects use cutpoint decision rules, how do cutpoints vary across treatments? How much variation is there

across individuals? How consistent is individual behavior with cutpoint decision rules? We document that there is some heterogeneity across subjects, but more importantly, the distribution of these cutpoint strategies varies systematically across treatments.

In order to estimate cutpoint 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 ϕ , 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 and condition. If there are multiple best fitting cutpoints, we take the average. Table 7 presents some cutpoint summary statistics.

Condition	M	median estimated cutpoint	% misclassified	empirical optimum
Simultaneous	.39	37.5	3.4	17
Simultaneous	.50	45.5	3.7	26
First mover	.39	42.5	2.1	15
First mover	.50	42.0	2.3	18
Second mover	.39	36.0	0.0	18
Second mover	.50	38.0	0.8	26

Table 7. Cutpoint summary statistics.

Several observations are immediate. First, very few decisions are misclassified. In each of the simultaneous treatments, 16 of 28 subjects are *perfectly* classified. In the sequential conditions, the number of perfectly classified subjects range from 23 to 28 of 28. The worst case of misclassification was one subject in the simultaneous, $M = .50$ treatment who has 5 misclassified observations. Thus, with rare exceptions, subjects use cutpoint strategies. In our setting, it establishes that subjects have a basic understanding of the game, even if they do not play the Nash equilibrium strategy. Our subjects follow the first step of the logic (best responses are cutpoints), but fail at the second one (equilibrium unraveling). Although interesting, this result is not the central point of the section. We document it to justify our approach of comparing the distributions of estimated cutpoints across experimental conditions.

The key questions concern whether the distribution of cutpoints varies systematically across treatments and, in the sequential games, between first and second movers. We also want to understand whether these systematic variations are consistent with the descriptive findings based on aggregate data described in the previous section. Table 7 reports the median estimated cutpoint across all subjects, and the percentage of misclassified decisions, by condition. The median cutpoints mirror the aggregate fight rates by treatment and condition, as reported in previous sections. Cutpoints are lower (more fighting) in $M = .39$ than in $M = .50$ treatments. They are lower for second movers than first movers for both the $M = .39$ and $M = .50$ treatments, and they are also lower for second movers than for players in the simultaneous condition for both the $M = .39$ and $M = .50$ treatments. There is no systematic difference between first movers and players in the simultaneous condition. Last, second movers have fewer classification errors (therefore, steeper response curves) than first movers.

As further evidence, we consider how the entire distribution of individual estimated cutpoints varies across treatments. Figure 2 displays the estimated cumulative frequency distribution of individual cutpoints used by our subjects for all the treatments, and broken down by position in the sequential treatment. The horizontal axis represents cutpoints, ranging from 0 to 100. The vertical axis indicates how many of the subjects in each treatment or position (out of 28) were using a cutpoint less than or equal to that number.

[FIGURE 2 HERE]

These distributions exhibit a wide range of estimated cutpoints, with few above 60 or below 20. In *all* treatments and conditions, there is heterogeneity that is both significant (one easily rejects the hypothesis that all subjects use the same cutpoint for any of these conditions), and substantial. The distributions are also different across treatments and conditions. Also noteworthy is that distributions are never concentrated around particular cutpoints (i.e., step distribution functions), but they are smoothly and uniformly increasing over the range.

Finally, one can compare the distribution of cutpoints used by players in the game to the cutpoint that would be optimal, given the actual frequencies of fighting in the experiment. These "empirically optimal" cutpoints are given in the last column of table 7. The optimal cutpoints are generally about one-half times the corresponding median estimated cutpoints. Many, but not all, players are "fooled" by this game, in the sense that they set

cutpoints that are too high. We find that 20% of the estimated cutpoints are within 5 units of strength of the optimal cutpoint, and these subjects are leaving essentially no money on the table. Of the remaining estimated cutpoints, 7% are below the optimal cutpoint by at least 5 strength units, and 73% are above the optimal cutpoint by at least 5 units.

4.4 Summary of descriptive analysis

The main findings of our analysis so far can be summarized as follows.

- Unconditional fight rates range from about 50% to 70%, depending on the treatment and condition, falling far short of the theoretical prediction of 100%.
- In all treatments, the ϕ -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 $M = .50$.
- The compromise payoff, M , affects behavior, with less fighting when the compromise payoff is higher.
- There are some differences between the sequential and simultaneous treatments. Most striking is that second movers behave differently than first movers: the former display more fighting and less erratic behavior (when $M = .39$) than the latter. For a theory to explain this pattern, it must predict that observing the behavior of the rival before making inferences and choices leads to systematic differences compared with just conditioning on a hypothetical event.
- We do not find consistent evidence of learning. A possible explanation is the insufficient feedback provided to players (only the rival's strength and 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 experience in that role when they subsequently play as first mover.⁸
- The vast majority of subjects use cutpoint strategies, with very few deviations. Across all treatments, over 97% of individual behavior is

⁸Recall 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.

consistent with cutpoint strategies. This shows their understanding, at least at an intuitive level, that the expected payoff differential between ϕ and ρ increases with s , and possibly an even deeper understanding that the best reply in this game is *always* a cutpoint strategy. It also justifies the analysis that focuses on estimated cutpoints.

- The distribution of these cutpoints varies by condition in ways that mirror the differences in the aggregate fight rates. The empirical distribution of individual cutpoints is smooth.

5 Alternative behavioral models

Obviously, the Nash equilibrium model is inconsistent with our data. In this section, we consider several alternative 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 that maps strength into a probability of fighting $q : [0, 1] \rightarrow [0, 1]$. 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 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.⁹ 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 .

We consider two categories of models. Models in the first category have their foundations in *cognitive limitations*, and they all have features that admit the possibility of observing weakly dominated strategies. We study three models within this class: quantal response equilibrium or QRE (McKelvey and Palfrey, 1995), cognitive hierarchy or CH (Camerer et al., 2004), and cursed equilibrium or CE (Eyster and Rabin, 2005).¹⁰ We also consider some variations that allow for heterogeneity or hybridization between models, such as the truncated quantal response equilibrium or TQRE (Camerer

⁹For the $M = .50$ game, at the last iteration, a player with the lowest strength, $s = .01$, is indifferent between ϕ and ρ , and therefore, there is an equilibrium with $s = .01$ types choosing ρ and all other types choosing ϕ . For the $M = .39$ game, the iteration continues all the way down, and the only equilibrium is $q^*(s) = 1$ for all s .

¹⁰Some preliminary findings about CH and CE are discussed in Wang (2006), with permission of the authors.

et al., 2006). These hybrid versions allow us to understand better how different models capture different features of the observed behavior. Models in the second category have their foundations on the existence of *systematic deviations from self-interest*: social preferences, fairness motives, reciprocity, or altruism. For reasons we discuss later, we estimate only models in the first category.

5.1 Quantal Response Equilibrium

Quantal response equilibrium applies stochastic choice theory to strategic games. It is motivated by the idea that a decision maker may take a sub-optimal 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 correspondence 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, λ , is a measure of responsiveness of choices to payoffs. That is:

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

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

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. 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 criterion 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. Last, focusing on cutpoint strategies does not seem too restrictive, given that the behavior of individuals is highly consistent with this type of play, as previously documented.

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 λ we solve for a fixed point in behavioral strategies. Denote by ϕ_λ^* such an equilibrium fixed point, and by $\phi_\lambda^*(s)$ the equilibrium probability of fighting given a strength s .

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

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

The expected utility of ρ conditional on the other player having chosen ρ is simply $V_\rho(s; \phi_\lambda^*) = M$, so the difference in the expected utility of ϕ and ρ 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, ϕ_λ^* 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. Therefore, 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 ρ . 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 ϕ if $s_i \geq c$ and chooses ρ 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 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:

$$U(\tilde{c}) = \int_{\tilde{c}}^1 s ds + \int_0^{\tilde{c}} \left[\int_0^s q(c) (cM + (s - c)) dc + \int_s^1 q(c) cM dc \right] ds \quad (2)$$

The first term is the probability of drawing a strength s above the cutpoint, in which case player 1 chooses ϕ 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 ρ . Then, if player 2's strength is lower, a compromise gives payoff M and a no-compromise gives payoff 1. If player 2's strength is higher, 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:

$$U_1(\tilde{c}) = \int_{\tilde{c}}^1 s_1 ds_1 + \int_0^{\tilde{c}} \left[\int_0^{s_1} q_2(c) (cM + (s_1 - c)) dc + \int_{s_1}^1 q_2(c) cM dc \right] ds_1 \quad (3)$$

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 ρ . 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 QRE-cut 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 ϕ rates. 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 logit QRE models by standard maximum likelihood techniques, i.e., finding the value of λ that maximizes likelihood of the observed frequencies of strategies. We estimate 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 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.

5.2 Cognitive Hierarchy

The cognitive hierarchy 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 . The CH solution to a game is uniquely determined by an assumption about how level 0 types behave (σ_0), and the distribution of levels of sophistication (p).¹¹ 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 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 then 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 $b^k(j) = p_j / \sum_{l=0}^{k-1} p_l$ for all $j \in \{1, \dots, k-1\}$.

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

¹¹The CH model is an extension of the original level- k model of Nagel (1995). See Stahl and Wilson (1995), Crawford and Iriberri (2005), Costa-Gomes and Crawford (2006), and Camerer et al. (2006) for further examples of estimation of level- k models, based on experimental data.

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 ϕ or ρ , 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 cognitive hierarchy model or CH-cut, 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 ϕ when their strength is high than when their strength is low. In our game, a level 0 type who randomizes over cutpoints has a probability of fighting 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 ρ , the posterior distribution of strength of a level 0 type conditional on choosing ρ is $f(s|\rho) = \frac{1-s}{\int_0^1(1-x)dx} = 2-2s$. Hence, the expected payoff of ϕ for a level 1 type with strength s and conditional on the other player being level 0 and choosing ρ is $\int_0^s(2-2x)dx = 2s-s^2$. Since the payoff of ρ 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 ϕ probabilities are monotone in s for all players.

We fit the Poisson specification of the CH and cutpoint CH models to the data set by finding the value of τ that maximizes likelihood of the observed aggregate frequencies of strategies, under the assumption that types are identically and independently distributed draws. We estimate the best-fitting values of τ by maximum likelihood for each of the four treatments, and report both constrained and unconstrained estimates.

5.3 Combining quantal response and strategic hierarchies (TQRE)

The predictions of the CH and CH-cut models differ from the QRE and QRE-cut models in two important ways. First, in CH models, 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. Neither "bunching" by layers of reasoning nor identical behavior in the simultaneous and sequential treatments are observed in the data.

An approach that combines quantal response and hierarchical thinking, called *Truncated Quantal Response Equilibrium (TQRE)*, is developed in Camerer et al. (2006). This model introduces a countable number of players' *skill levels*, $\lambda_0, \lambda_1, \dots, \lambda_k, \dots$. The distribution of skill levels in the population is given by $p_0, p_1, \dots, p_k, \dots$. A player with skill level k chooses stochastically with a logit quantal response function with precision λ_k . TQRE assumes truncated rational expectations in a similar manner to CH: a player with precision λ_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 spaced $\lambda_k = \gamma k$. Thus, it is a two parameter model with Poisson parameter, τ , and a spacing parameter, γ .

The TQRE model has two effects. It smooths 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, so we estimate both versions.

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 α -cursed equilibrium (CE_α) all players are α -cursed. However, players believe that opponents are α -cursed with probability $(1 - \alpha)$ and they believe that actions of opponents are independent of their information with probability α . 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 CE_α 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 ϕ , conditional on the opponent choosing ρ is given by:

$$\begin{aligned} V_\phi^\alpha(s_i) &= \alpha \Pr\{s_j < s_i\} + (1 - \alpha) \Pr\{s_j < s_i \mid a_j = \rho, s_\alpha^*(M)\} \\ &= \alpha s_i + (1 - \alpha) \min\left\{1, \frac{s_i}{s_\alpha^*(M)}\right\} \end{aligned}$$

A player with strength equal to the equilibrium cutpoint must be indifferent between ϕ and ρ . Formally, $V_\phi^\alpha(s_\alpha^*(M)) = V_\rho^\alpha(s_\alpha^*(M))$. Therefore:¹²

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

A difficulty with CE_α is that it cannot be fit to the data due to a zero-likelihood problem: for each α 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 CE_α .¹³ In the simultaneous move game, a (symmetric) α -QRE is a behavior strategy, or a set of probabilities of choosing ϕ , one for each value of $s \in [0, 1]$. We denote such a strategy evaluated at a specific strength value by $\phi(s)$. Given λ and α we denote by α -QRE the behavior strategy $\phi_{\lambda\alpha}^*$. If player j is using $\phi_{\lambda\alpha}^*$ and player i is α -cursed, then i 's expected payoff from choosing ϕ when $s_i = s$ is given by:

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

And the expected payoff from choosing ρ is:

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

¹²In 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: $s_1^*(M) = M$.

¹³Note that player heterogeneity with respect to α would 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 ρ .

Using the logit specification for the quantal response function, we then apply logit choice probabilities to the difference in the expected payoff from ϕ and ρ 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 α -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 α .

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 ϕ and ρ for the first mover are the same as in the simultaneous move game, so we have:

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

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

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

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 Δ_2 is similar to the RHS of Δ_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 α -QRE model to the data set by finding the values of λ and α that maximize likelihood of the observed frequencies of strategies. As for the previous models, we report both constrained and unconstrained estimates.

5.5 Models of pro-social behavior

We also considered an alternative class of models which are not based on cognitive limitation but, instead, are founded on social preferences. There are a number of candidates from this growing family of models. We consider three. One is the *fairness* model by Fehr and Schmidt (1999). In that model, the utility of individual i when he gets payoff x_i and individual j gets payoff x_j is:

$$U_i(x_i, x_j) = x_i - \alpha \max\{x_j - x_i, 0\} - \beta \max\{x_i - x_j, 0\}$$

where $\beta \leq \alpha$ and $0 \leq \beta < 1$. For our game, the model implies that the utility payoff to each agent for winning, compromising and losing are $1 - \beta$, M , and $-\alpha$, respectively. This implies that if fairness considerations are sufficiently strong ($\beta \geq 1 - M$), the equilibrium unravelling goes the opposite direction, and *all* agents always play ρ , regardless of their strength. Otherwise, agents want to set a lower cutpoint than their rival, and we are back to the Nash equilibrium prediction where all agents always play ϕ . Our subjects do not exhibit such extreme "boundary" behavior, neither individually nor in the aggregate. Thus, one would need to add other parameters or assumptions in order for this model to explain the choices of our subjects.

A second model is *altruism*, where a player's utility is the weighted average of his own payoff and the other agent's payoff:

$$U_i(x_i, x_j) = \gamma x_i + (1 - \gamma) x_j$$

This model runs into the same problem as the previous one. Each player’s payoff of winning, compromising and losing are γ , M , and $1-\gamma$, respectively. Therefore, all players should either always fight or always retreat. Given estimates of γ from other experiments ($\gamma > .5$), the model predicts that players should always fight if $M \leq .50$.

Third, models based on *reciprocity* are also prominent in the social preferences literature. These models provide an explanation for the behavior commonly observed in the trust game. In our setting, they suggest that second movers should fight less than first movers, as they are “returning the favor” of compromising. However, we find *the opposite*: observing ρ significantly *decreases* rather than increases the willingness to reciprocate by responding also with ρ .

Overall, these leading models of social preferences described above – fairness, altruism, and reciprocity – fail to explain the basic patterns we observe in the data. Therefore, we do not estimate them. In fact, there are at least two additional reasons why models of pro-social behavior are unsuitable to account for the choices of subjects in this particular game. First, each individual plays the game many times (40), anonymously against a pool of opponents, with the roles of players, and the strengths randomly assigned in each match. Given this design, applying models of social preferences to behavior in isolated games is questionable a priori. In fact, in one of the designs ($M = .50$), subjects play a constant sum game against changing opponents. Hence, deviations from optimal best replies to “the field” will necessarily, over the course of the 40 matches, give a player a total expected payoff below the average payoff of the other subjects. Second and related to the above argument, in this game *selfish play leads to ex ante fair and efficient allocations*. Indeed, for the case of $M = .39$, myopic “fair” behavior (everyone retreating every time) leads to long run inefficient outcomes, and no long run improvement in the equality of payoffs. Subjects would be leaving over 20% of potential group earnings on the table, with virtually zero gain in equality of outcomes.

5.6 Model estimates

In this section we estimate the QRE, CH, TQRE and CE models, and explore the stability of the estimated parameters across the different treatments, and compare the ability of these models to capture the basic features of the data, identified in the previous section. We report the estimates in Table 8 at different levels of aggregation: for the treatments separately; pooling

across the M -treatments; and pooling across all treatments. For the QRE, CH and TQRE models, we considered both the behavioral strategy version and the cutpoint version. In all three models, the cutpoint version fits 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 and discuss only the results for the cutpoint versions of these models.

	N	QRE		CH		TQRE			α -QRE		
		λ	$-\ln L$	τ	$-\ln L$	γ	τ	$-\ln L$	λ	α	$-\ln L$
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 8. Model estimates.

Figure 3 displays the empirical and fitted fighting probabilities as a function of strength. Fitted choice frequencies in the figure are based on out-of sample parameter estimates. Specifically, the displayed curves for the sequential data are constructed using the parameter estimates obtained from the pooled simultaneous data, and vice versa. All these models capture the upward sloping empirical frequency of ϕ . All exhibit low ϕ rates for low strengths and high ϕ rates for high strengths.

[FIGURE 3 HERE]

There is some variation in fit across the different models. The better fitting models all converge to 0% for low strengths and to 100% for high strengths. The α -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. The CH and QRE models fit the data similarly in terms of log likelihood, but there are some important differences in the predicted fight curves. As one can see from figure 3, the CH model predicts somewhat better at strengths below 20%, but QRE generally fits better than CH elsewhere. This can be attributed to two important differences in the models: (1) QRE predicts that second movers will have different

(and sharper) response functions than first movers, which is a feature of the data not captured by CH; (2) QRE generates a smooth fight curve, while CH predicts clustering of cutpoints, leading to jumps in the fight curve corresponding to different levels of sophistication.

The TQRE model does not provide a substantial improvement over QRE or CH. In fact, the fitted ϕ -rate for TQRE and QRE are very similar. They both share the problem of overestimating the fighting rates for subjects with low strength. The α -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 α , the players' degree of cursedness. Adding quantal response, produces a nice logit function of the fighting probability, that crosses .50 at $s \approx .40$, varying slightly with M and position, consistent with the data. Furthermore, quantal response also introduces a steeper ϕ 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 $M = .39$ and the $M = .50$ treatments, with most models fitting the data from the $M = .39$ treatment better, reflecting the steeper empirical ϕ curves in the $M = .39$ data. There is virtually no difference in either the fit or the actual parameter estimates for the sequential and simultaneous treatments. The α -QRE pooled estimates of λ and α are not significantly different between the two treatments, even at the 5% level, and the fit is identical ($\log \text{Likelihood} / N = -.318$ in both cases).

5.7 Summary of estimation results

The main findings about the estimated models summarized as follows.

- All four models capture the most basic qualitative properties of behavior (none of which is consistent with Nash equilibrium): fight rates are high, increasing in s , and decreasing in M . However, each model captures different specific features of the data.
- Estimates are similar across M treatments, and little power is lost by pooling treatments. This means that the results are not due to "tuning the parameters to fit the data". In fact, the out-of-sample parameters from $M = .39$ provide virtually identical predictions for $M = .50$ behavior as the within sample estimates.
- Only QRE and the models hybridized with QRE capture the fact that

first movers behave differently from second movers. In particular, the ϕ function is steeper for second movers in those models.

- The cutpoint versions of CH and QRE describe behavior better than the behavior strategy versions. This is consistent with our findings at the individual level which indicate that over 95% of choices follow pure cutpoint strategies.
- TQRE provides an almost identical fit as QRE, suggesting that, in this game, the addition of hierarchical thinking to quantal response does not have a substantial impact. This is also consistent with the fact that we do not find individual cutpoints clustered around 3 or 4 strength values, as would be predicted by CH and other levels-of-sophistication models.
- The α -QRE model fits the data best. The estimates of α are significantly greater than 0 and significantly less than 1. They are virtually identical for both the sequential and simultaneous games, suggesting that the 2-parameter model is not overfitting the data.

6 Conclusions

The compromise game is obviously 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 best responses. However, they have problems figuring out the full logic of the unravelling argument.

The paper has considered several cognitive explanations for the surprising behavior observed in these games of incomplete information. In a future research, it might be interesting to explore more general models. One candidate would be the "analogy-based expectation equilibrium" developed by Jehiel and Koessler (2006), which can be seen as a generalized version of cursed equilibrium. A second direction would be to explicitly allow for heterogeneity. While the CH model is suggestive of heterogeneity, the attempts here and 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 estimation of CH models to allow for fixed types. However, for our data, it seems unlikely to go very far because we do not observe clusters of behavior that might correspond

to types – in contrast to Nagel’s (1995) guessing game, for example. The QRE and α -QRE models could also be extended to allow for heterogeneity with respect to λ and α , also with fixed types. This would undoubtedly lead to better fits in terms of log likelihood, but it is hard to imagine any new insights emerging from such an exercise.

One of the most interesting findings is that the order of moves affects choices. In our game, a player’s action is relevant only if the rival chooses ρ . Thus, first, second and simultaneous movers should all condition their strategy on that event. By contrast, the data shows that players who observe ρ 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). In sum, *hypothetical conditioning on events seems to produce different behavior than observational conditioning on events*. While we found one general explanation for this phenomenon (QRE), a search for other parsimonious formal models that imply different behavior between our first and second movers might add to our understanding. Such a search seems a worthy project for future research. It could also help settle questions well beyond the scope of our compromise game, such as the effect of using the strategy method in experimental games, and the differences in behavior between strategically equivalent games (extensive vs. strategic form). Naturally, this has implications for many strategic settings of significant applied interest, including common value auctions and voting behavior, where optimal choice requires bidders to condition on winning and voters to condition on being pivotal.

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, 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*, 57: 123-54
- [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. Iriberry (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] Holt, C. and R. Sherman (1994), "The Loser's Curse," *American Economic Review*, 84(3), 642-652.
- [12] Jehiel, P. and F. Koessler (2006), "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *Working Paper*, PSE.
- [13] Kagel, J and D. Levin (2002) *Common Value Auctions and the Winner's Curse*, Princeton University Press: Princeton, NJ.

- [14] 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.
- [15] McKelvey, R. and T. Palfrey (1995) "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- [16] McKelvey, R. and T. Palfrey (1998) "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1, 9-41.
- [17] Milgrom, P. and N. Stokey (1982) "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17-27.
- [18] Morris, S (1994) "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327-1347.
- [19] Myerson, R.B. and M.A. Satterthwaite (1983) "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, 29(2), 265-281.
- [20] Nagel, R. (1995) "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5), 1313-1326.
- [21] Palfrey, T. and J. Prisbrey (1996) "Altruism, Reputation, and Noise in Linear Public Goods Experiments," *Journal of Public Economics*, 61, 409-27.
- [22] 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.
- [23] 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.
- [24] Sovic, Y. (2004) "Strength of Dominance and Depths of Reasoning," *Working Paper*, University of Oslo.
- [25] Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10(July), 218-54.
- [26] Wang, S. (2006), "Misunderstanding the Other," *Mimeo*, 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. Empirical and Fitted Fight Rates.