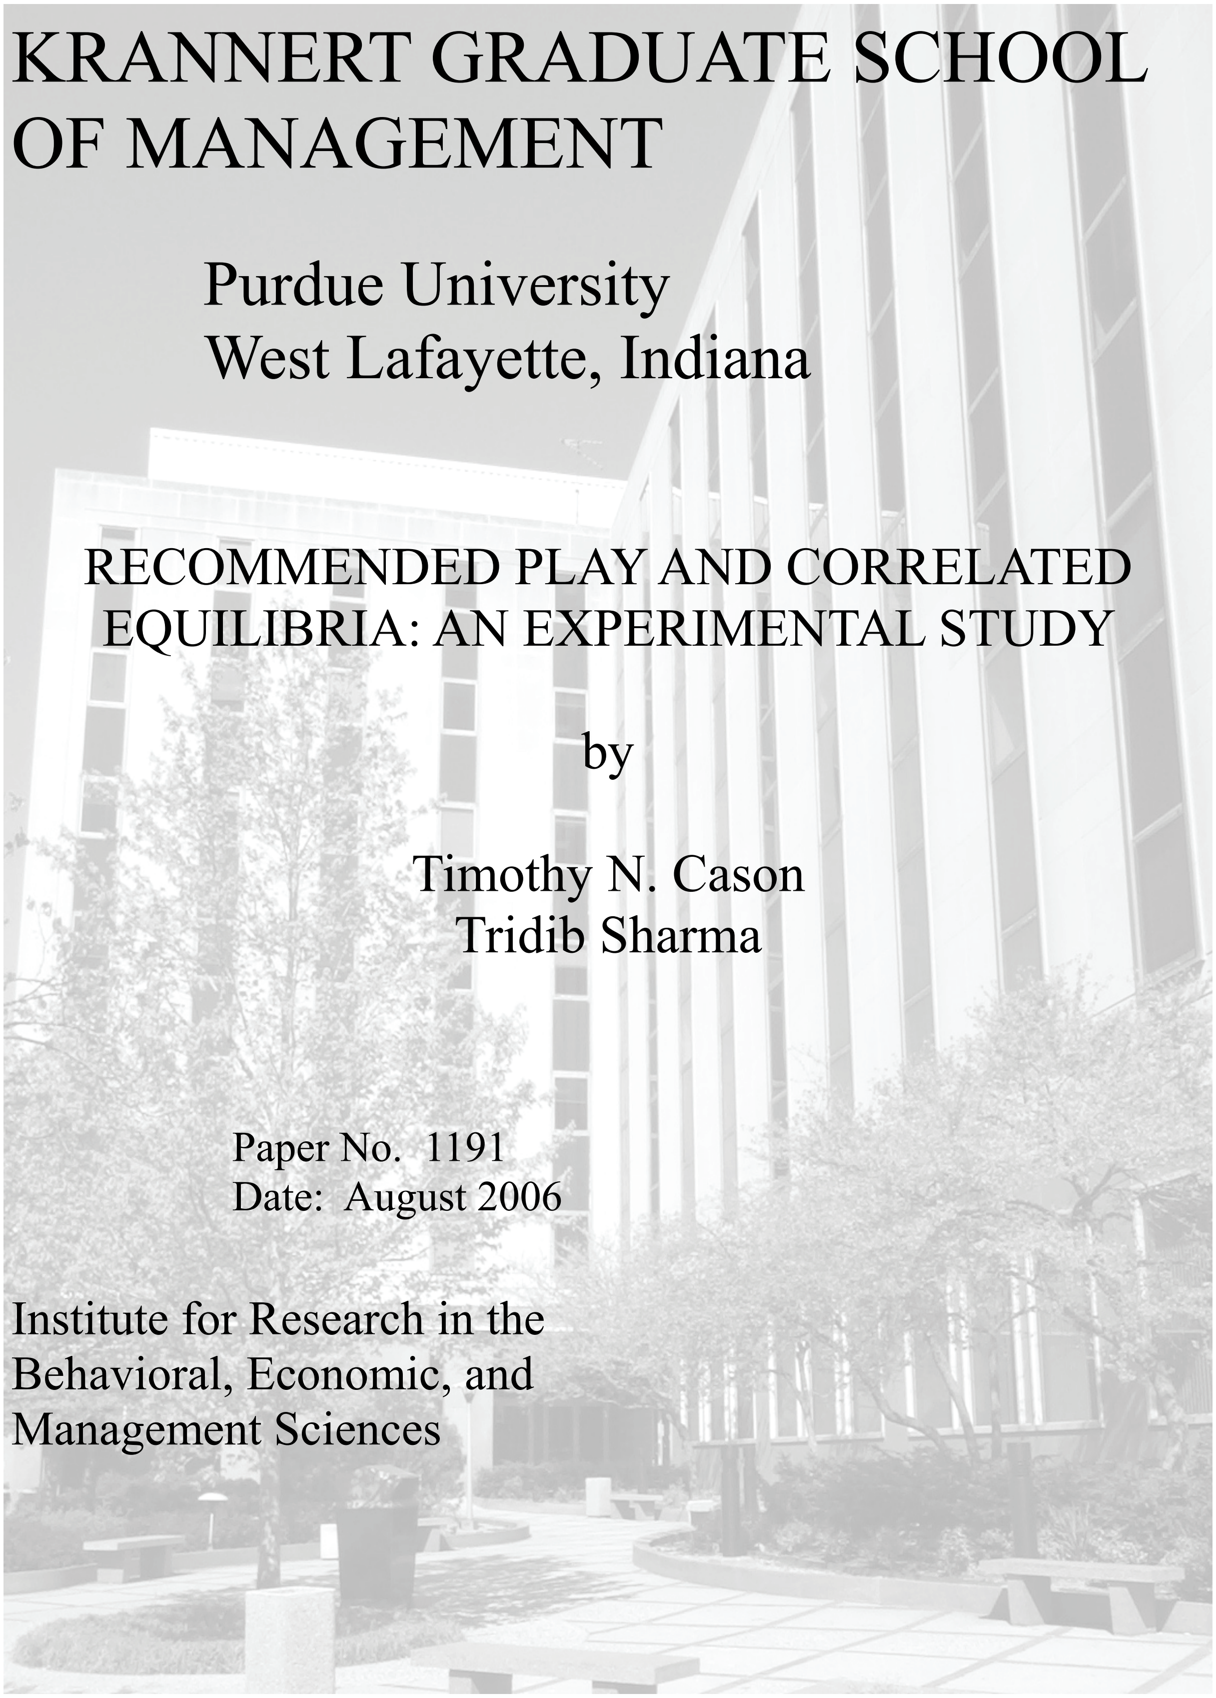

KRANNERT GRADUATE SCHOOL OF MANAGEMENT

Purdue University
West Lafayette, Indiana

RECOMMENDED PLAY AND CORRELATED EQUILIBRIA: AN EXPERIMENTAL STUDY

by

Timothy N. Cason
Tridib Sharma

Paper No. 1191
Date: August 2006

Institute for Research in the
Behavioral, Economic, and
Management Sciences

Recommended Play and Correlated Equilibria: An Experimental Study*

Timothy N. Cason
Purdue University
100 S. Grant Street
West Lafayette, IN 47907-2076, USA

Tridib Sharma
Instituto Tecnológico Autónomo de México
Ave. Santa Teresa 930
Mexico City 10700, Mexico

August 12, 2006

Abstract

This study reports a laboratory experiment wherein subjects play a hawk-dove game. We try to implement a correlated equilibrium with payoffs outside the convex hull of Nash equilibrium payoffs by privately recommending play. We find that subjects are reluctant to follow certain recommendations. We are able to implement this correlated equilibrium, however, when subjects play against robots that always follow recommendations, including in a control treatment in which human subjects receive the robot "earnings." This indicates that the lack of mutual knowledge of conjectures, rather than social preferences, explains subjects' failure to play the suggested correlated equilibrium when facing other human players.

JEL classification: C72

Keywords: game theory; experiments; coordination; common knowledge

*We are grateful for financial support provided by the Purdue University Faculty Scholar program and the Asociación Méxicana de Cultura, as well as for the valuable research assistance provided by Shakun Datta and Marikah Mancini. We received helpful comments from Shurojit Chatterji, David Cooper, Arthur Schram, Ricard Torres, an anonymous referee, and from conference and seminar participants at Royal Holloway, the University of Amsterdam, Purdue University, the Economic Science Association and the Society for the Advancement of Economic Theory.

1 Introduction

Consider a normal form game with multiple Nash equilibria, and suppose we observe what players play. We may then determine whether or not play corresponds to some particular equilibrium with desirable characteristics. If not, then can we induce players to choose this equilibrium? In this paper we report results from experiments which try to execute a switch to a "better" (in a sense defined later) equilibrium. We find the posed question relevant because games often have equilibria with undesirable characteristics. Even when games are designed such that an equilibrium of the game satisfies some pre-set objective, it is not always possible to achieve a unique equilibrium. For example, under moral hazard in teams the designed wage structure may lead to multiple Nash equilibria (Holmstrom, 1982; Mookherjee, 1984). In such situations, one might want to design games that explicitly incorporate coordination mechanisms. Under moral hazard in teams, once the wage schedules are accepted one can employ a manager who simply recommends to each worker the amount of effort that each of them should input (Sharma and Torres, 2004). If the recommendations are incentive compatible then theory tells us that they should be followed and therefore the equilibrium that is desired should actually result. Is this really so, or is there a slip between the cup and the lip? The goal of our paper is to investigate this question.

The experiments reported in this paper do not exhaust all the possible means to induce coordination. Instead, we adopt a simple procedure. We instruct players to play their respective strategies of a selected Nash equilibrium to implement a correlated equilibrium. Simple as the procedure is, it is not without problems.[1] First, our experiments provide subjects with monetary payments. We do not observe the players' utility (payoff) over

[1] In theoretical terms the payoffs, rationality and conjectures may not be common knowledge amongst the players and the experimenter (see for example Aumann, 1998, and Gul, 1998).

these payments. Hence, we cannot be certain about the set of equilibria in the actual game that subjects play. Second, even if we were to know the payoffs, the designed incentive compatible instructions may still fail to induce common knowledge of players' conjectures over their opponent's play. This may happen because payoffs or rationality are not common knowledge amongst players (for more on this see Aumann and Brandenburger, 1995). In this paper we try to disentangle these two potential problems.

To fix matters, consider the following symmetric hawk-dove game that is used in our experiments (we shall motivate our choice of this game in the next section):

	Left	Right
Up	3, 3	48, 9
Down	9, 48	39, 39

The values in each cell are monetary transfers from the experimenter to the subjects. If we assume that these transfers are equivalent to utilities, then the game has three Nash equilibria. In addition to the {Down, Left} and {Up, Right} pure strategy Nash equilibria, it has a mixed strategy Nash equilibrium in which players choose Up and Left with probability 0.6. Mutual, and not necessarily common, knowledge of payoffs, rationality and conjectures could result in one of these Nash equilibria (Aumann and Brandenburger, 1995).

Our assumption on payoffs could, however, be wrong. Suppose Row's utility is equal to his monetary transfer less Column's monetary transfer, a very simple form of rivalistic social preferences.[2] Then playing Up is a dominant strategy for him and he will not play Down if instructed to do so. Now suppose our assumption that money transfers are equivalent to utilities is correct. But also suppose that Row believes with probability one that Column is rational and Column's payoff is equal to Column's payoff minus Row's payoff. Thus payoffs

[2] For recent influential models of social preferences, see Fehr and Schmidt (1999) or Bolton and Ockenfels (2000).

are not mutual knowledge. Then Row believes that Left is a dominant strategy for Column and Column will always reject a recommendation to play right. Hence Row will always ignore the instruction to play Up. Thus we see that a player may reject instructions for two reasons. First, given his payoff it is optimal for the player to not follow instructions. This is an example where our assumption on payoffs is wrong. Second, our assumption on payoffs is correct, but the player does not follow his recommendation because he believes that his opponent will not follow her recommendation.

We start our exercise by observing randomly matched subjects play the game. We assume that payoffs are equivalent to monetary transfers. Our observations cannot reject the hypothesis that the subjects are playing the mixed strategy equilibrium where Up and Left are chosen with probability 0.6. Given this finding, we proceed to implement a different Nash equilibrium.

To induce a different Nash equilibrium we recommend an action to each player. Recommending actions actually transforms the original game into a signalling game where each player receives a signal (which does not affect payoffs) prior to the choice of strategies. The Nash equilibrium of the game with signals is then a correlated equilibrium of the original game. The set of all consistent signal structures then generate the set of all correlated equilibria. The set of correlated equilibria includes the Nash equilibria of the original game as these can be seen as resulting from independent signals (Hart and Mas Colell, 2000). Thus, by appropriately choosing recommendations (signals) we can induce players to coordinate on a given correlated equilibrium. Assuming that subjects' payoffs are equal to their monetary transfers we try to induce a correlated equilibrium "close" to the symmetric correlated equilibrium which maximizes joint payoffs. This equilibrium puts probabilities 0.375, 0.375, 0.25 on {Up, Right}, {Down, Left} and {Down, Right}.

Suppose the experimenter wants to implement this "optimal" symmetric correlated equilibrium. He publicly announces a loaded dice with three faces {Up, Right}, {Down, Left} and {Down, Right} that occur with probabilities 0.375, 0.375, 0.25. The subjects do not get to see the outcome of dice rolled by the experimenter. The experimenter sees the outcome and privately recommends the respective strategies to the players. Suppose the row player is given the recommendation to play Down. Assuming he is able to correctly update probabilities using Bayes' rule, then he would know that with probabilities $\frac{.375}{.625}$ and $\frac{.25}{.625}$ Column has been asked to play Left and Right. (It is well known, however, that laboratory subjects frequently fail to use Bayes' rule.) If Row were to also know that Column follows recommendations then this knowledge is sufficient for him to make an optimal choice. In particular, if his payoff is equal his monetary transfer, then optimization would induce him to follow the recommendation and choose Down.

When we run the experiments with recommendations we find that players frequently reject recommendations. Why does this occur? It could simply be that our payoff specifications are wrong. Or it could be that our specifications are correct but players do not believe that their opponents will follow recommendations.

To investigate the problem we first run a treatment where each subject plays a robot. It is announced that robots will always follow their recommended strategy. We observe in this treatment that subjects almost always follow recommendations. This finding suggests that if opponents were known to follow their recommendations then it is optimal for the subject to also follow recommendations. Since optimization is done by subjects over their payoffs, we are unable to reject the hypothesis that we had specified payoffs correctly. However, payoffs could still be misspecified in the game with human subjects. This is because playing against robots should substantially reduce or even eliminate concerns about social preferences. To check whether social preferences cause the problem when subjects face human opponents, we repeated the robot treatments with a variation. The monetary amounts "earned" by the robot are transferred to a non playing human subject. Again we observed that recommendations were almost always followed. This indicates that recommendations are not followed when subjects face human opponents primarily because they believe that their opponents will not follow recommendations.

We also determined if experience could induce subjects to believe that their opponents would follow recommendations. We brought back subjects who followed recommendations in the robot opponent treatments to play each other. This group rejected recommendations as often as another experienced group who always played against human subjects.

In answering the question that we started with, our results indicate that communication through instructions is not sufficient to induce players to switch to *any* chosen equilibrium. It seems that this is not because players do not want to switch to the recommended equilibrium, but because they believe that their opponents will not switch. We shall provide further thoughts on this in the concluding section.

2 The Game

As mentioned above we study behavior in a simple, symmetric hawk-dove game. A feature of this game which is suitable for our purpose is that it has multiple (three) Nash equilibria. Furthermore, it has correlated equilibrium outcomes with higher joint payoffs that do not lie in the convex hull of Nash equilibrium outcomes. Thus, recommendations through an impartial mediator could potentially induce these superior outcomes.

Our game could occur naturally as an outcome of a design problem. To see this we construct a simple example following Sharma and Torres (2004). Consider a situation where two workers R and C have to jointly put in unobservable effort to produce output. Effort can be high or low, i.e. e_h or e_l with $e_h > e_l$. If both workers put in high effort then the value of the output is 156; if both put in low effort it is 6 and if one puts in low effort and other puts in high effort it is 96. Call these states H, L and M respectively. The workers have to design a wage schedule contingent on the value of the product, i.e. the state. For credibility, the sum of the wages should equal the value of the product in each state (Holmstrom, 1982). Given a state, constrain wages to be equal amongst workers for reasons associated with fairness or the unobservability of effort. It follows immediately that the wage schedule $\{w_L, w_M, w_H\} = \{3, 48, 78\}$. Let each worker have utility $w - e$, where w is the wage received and e is the level of effort put in. Let $e_h = 39$ and $e_l = 0$ for both workers. Then the game induced by the designed wage schedule is:

	e_l (Left)	e_h (Right)
e_l (Up)	3, 3	48, 9
e_h (Down)	9, 48	39, 39

This is exactly the game depicted in the introduction. In addition to the {Down, Left} and {Up, Right} pure strategy Nash equilibria, this stage game has a mixed strategy Nash equilibrium in which players choose Up or Left with probability 0.6.

If workers can communicate, then is it possible for them to coordinate their actions (i.e. play a correlated equilibrium)? Figure 1 shows the convex hull of the Nash equilibrium payoffs, as well as the set of correlated equilibrium payoffs. The set of correlated equilibrium payoffs can be characterized by assigning weights $(a_{11}, a_{12}, a_{21}, a_{22})$ (or a distribution over the four pure strategy profiles {Up, Left}, {Up, Right}, {Down, Left} and {Down, Right}) such that the following inequalities hold:

$$\frac{a_{11}}{a_{11}+a_{12}}3 + \frac{a_{12}}{a_{11}+a_{12}}48 \geq \frac{a_{11}}{a_{11}+a_{12}}9 + \frac{a_{12}}{a_{11}+a_{12}}39$$
$$\frac{a_{21}}{a_{21}+a_{22}}9 + \frac{a_{22}}{a_{21}+a_{22}}39 \geq \frac{a_{21}}{a_{21}+a_{22}}3 + \frac{a_{22}}{a_{21}+a_{22}}48$$

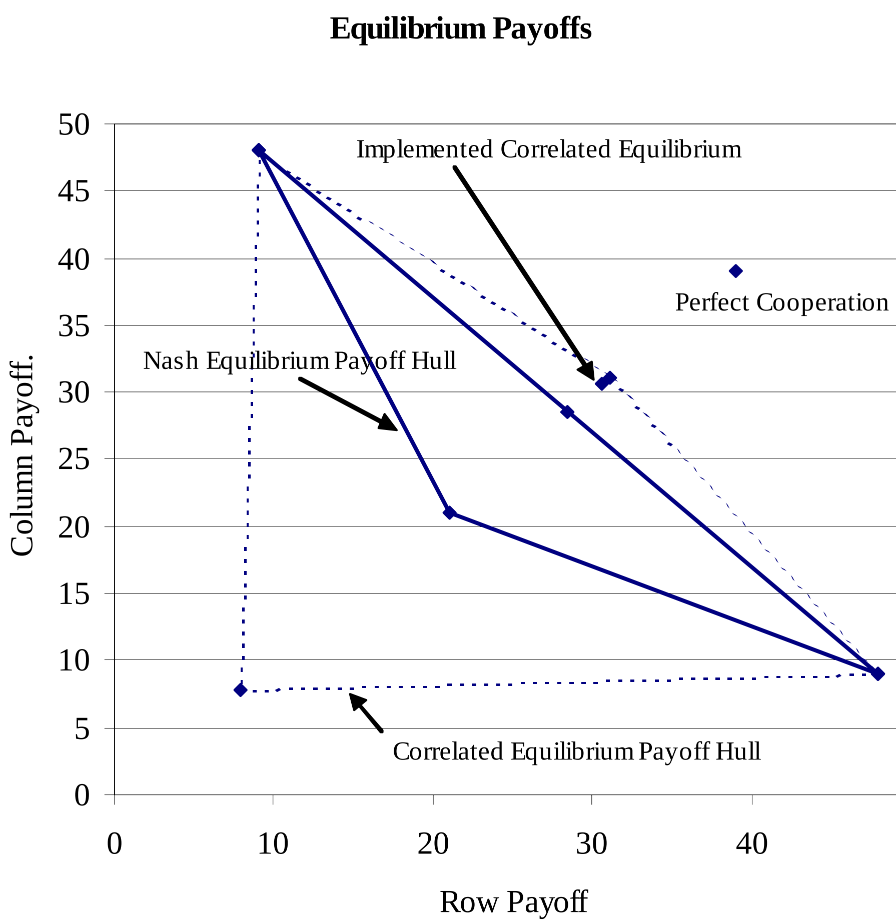
Equilibrium Payoffs
Implemented Correlated Equilibrium
Perfect Cooperation
Nash Equilibrium Payoff Hull
Correlated Equilibrium Payoff Hull
Column Payoff.
Row Payoff
0
5
10
15
20
25
30
35
40
45
50
0
10
20
30
40
50

Figure 1:

$$\frac{a_{11}}{a_{11}+a_{21}}3+\frac{a_{21}}{a_{11}+a_{21}}48 \geq \frac{a_{11}}{a_{11}+a_{21}}9+\frac{a_{21}}{a_{11}+a_{21}}39$$
$$\frac{a_{12}}{a_{12}+a_{22}}9+\frac{a_{22}}{a_{12}+a_{22}}39 \geq \frac{a_{12}}{a_{12}+a_{22}}3+\frac{a_{22}}{a_{12}+a_{22}}48$$

The idea here can be understood as follows. The workers employ an impartial mediator or manager (whose payoff, normalized to zero, is the same irrespective of the outcome). The manager recommends strategies for both players according to the commonly known distribution $(a_{11},\ a_{12}, a_{21}, a_{22})$. Given their knowledge, the distribution should be such that it is in the players' incentive to play the strategies that they are instructed to play. The first two inequalities are the incentive compatibility constraints for the row player to play Up and Down respectively. The second two inequalities for the column player to play Left and Right respectively. The set of all convex sets generated by these inequalities depict the set of all correlated equilibria payoffs. In particular, the convex hull of Nash equilibria payoffs lie strictly within this set.

Perfect cooperation payoffs (39, 39) lie outside the set of correlated equilibrium payoffs. Note that many points in the convex hull of correlated equilibrium payoffs are inferior to those in the Nash equilibrium. One of the key features that led us to use this game, however, is that there exist correlated equilibria that provide joint payoffs greater than those generated by any Nash equilibrium. For example, the highest symmetric payoffs that can be implemented as a correlated equilibrium provide a payoff of 31.125 to each player. Players are, however, indifferent between their action choices for recommendations in this equilibrium. So in order to make choices unique best responses for equilibrium beliefs we attempted to implement a correlated equilibrium strictly inside the hull of correlated equilibrium payoffs as shown in Figure 1. In this correlated equilibrium the costly {Up, Left} outcome is never realized. The pure strategy Nash equilibria are each played 40 percent of the time, and the cooperative outcome {Down, Right} is realized 20 percent of the time. In this equilibrium each player earns an expected payoff of 30.6. No Nash equilibrium of the stage game can generate these payoffs.

3 Experimental Design

Table 1 summarizes the experimental design and the allocation of human subjects to treatment conditions. All treatments with human opponents randomly rematched pairs of subjects each period to minimize the potential impact of reciprocal concerns. Treatment A serves as a baseline, in which subjects played the game with no recommendations. All other treatments employed play recommendations in an attempt to implement the correlated equilibrium described above. Theoretically, any correlating device could suffice to implement the equilibrium. Subtle devices such as blinking colored squares on subjects' computer screens to implement coordination do not work well unless subjects first experience real shocks to condition their beliefs to support coordination (Marimon et al., 1993). Therefore, we chose to be very explicit in our recommendations and the correlation process that leads to the recommendations. Each period we made a recommendation privately to each subject (e.g., "we recommend you choose left"). Following a pilot session in which some subjects appeared to ignore recommendations, for the main study we even added an explicit (and exhaustive) explanation of how subjects could update their beliefs following any recommendation, and how it was in their best interest to follow their recommendation if they believed that their opponent would also follow her recommendation. Our goal was to provide the most favorable and rich information conditions to make the use of signals transparent and valuable.[3]

[3] The instructions are available at www.krannert.purdue.edu/faculty/cason/papers/corr-inst.pdf. When referring to the subject's counterpart in the game, the instructions avoid using competitive framing; that is, we did not use the term "your opponent" unlike the presentation in this paper. The instructions also do not use obviously cooperative framing such as "your playing partner." Instead, they use more neutral (but also more clumsy) framing such as "the participant you are paired with."

Table 1: Allocation of Subjects to Treatments

Index	Treatment	Experience	Subjects	Sessions
A	Human opponents, no recommendations	None	52 subjects	4 sessions
B	Human opponents, with recommendations	None	80 subjects	6 sessions
C	Robot opponents, with recommendations, no one receives robot earnings	None	31 subjects	2 sessions
D	Robot opponents, with recommendations, humans receive robot earnings	None	26 subjects	2 sessions
E	Human opponents, with recommendations	Experienced in B	24 subjects	2 sessions
F	Human opponents, with recommendations	Experienced in C	24 subjects	2 sessions

Notes: Human opponents changed randomly each period, with typically 12 or 14 subjects participating in each session. Each session included 60 periods.

Several previous experiments have considered correlated equilibria or recommended actions in attempts to implement specific outcomes. Moreno and Wooders (1998) observe outcomes consistent with a correlated equilibrium in a three-player matching pennies game. Players achieve coordination through opportunities for nonbinding preplay communication. Other experiments have used experimenter instructions to implement specific equilibria in two-person coordination games (Van Huyck et al., 1992) or in four-person voluntary contributions games (Seely et al., 2005), and have found considerable difficulty in inducing subjects to follow recommendations. Brandts and MacLeod (1995) use experimenter recommendations to implement various types of Nash equilibria (e.g., dominant strategy, imperfect and subgame perfect). They also find that subjects do not blindly follow recommendations, particularly those corresponding to an imperfect equilibrium.

We also find that subjects frequently do not follow recommendations. To gather information about the source of this behavior, in all of our treatments we elicited beliefs from subjects regarding their expectations of their opponent's upcoming actions. They were paid based on the accuracy of their beliefs using a proper scoring rule (e.g., Nyarko and Schotter, 2002; Ruström and Wilcox, 2004).[4] These beliefs help provide insight into the learning process of subjects over time, as well as information about how expectations adjust after subjects receive play recommendations.[5] Subjects submitted their beliefs at the same time they made their action choice, after they received their action recommendation. Subjects' per-period earnings from their prediction accuracy varied between 0 and 10 experimental francs, so usually subjects earned more from playing the game than from making predictions.

In the robot opponents treatments C and D subjects knew that their counterparts were automated, and the instructions also emphasized that the robot opponents were programmed to always follow their play recommendation without error. This introduces knowledge over opponents following recommendations. But introducing robot opponents in treatment C also substantially reduces the possible influence of social preferences, such as when subjects care about the distribution of earnings between themselves and another person. For example, a row player may more strongly dislike her payoff in the {Up, Right} cell because her opponent earns 48 while she earns only 9. Such social preferences are very unlikely to affect subjects' choices in treatment C because the opponent earnings are not distributed to another human subject. Houser and Kurzban (2002) similarly use computerized opponents to limit the influence of social preferences. The problem that arises, however, is that any behavioral

[4] In particular, the payoff for a player reporting an Up prediction of r_{Up}, for example, when her opponent actually played Up, is $\pi_{Up} = 10 - 10(1 - r_{Up})^2$.

[5] As Ruström and Wilcox (2004) document, eliciting beliefs may also focus players' attention on their opponent's play and promote belief-learning. In theory, risk-averse or risk-seeking subjects might bias their reported belief to maximize expected utility, and a subject who integrates their belief and game choice decisions might also report biased beliefs. Evidence reported by Offerman et al. (1996), which includes belief elicitation from "spectators" not playing the game, indicates that these concerns are not empirically significant.

differences between treatments B and C could occur either because of differences in beliefs over opponents following recommendations or because of the elimination of social preference considerations. This is a classic problem of an experimental treatment confound.

The intermediate treatment D solves this confound problem. Like treatment C, subjects play against robots who always follow recommendations. Like treatment B, another human subject receives the earnings of the subject's opponent. Therefore, if subjects care about the distribution of earnings (as in well-known models such as Fehr and Schmidt, 1999, or Bolton and Ockenfels, 2000) then these social preferences will have a similar influence on choices in treatments B and D. In the treatment D sessions half of the subjects played against robots and their play determined their own earnings. The other half of the subjects simply received these robot earnings and they could not affect anyone's payoffs. These "bystander" subjects also played the same game 60 times (with recommendations) while they waited for the other subjects to complete their decisions, although the instructions emphasized that these decisions were irrelevant for anyone's payoffs. (We do not analyze these hypothetical choice data in this paper.) Subjects who did determine earnings were informed that the human who received their opponent's earnings changed randomly each round, as in all human opponent player treatments.

We report data from the 237 subjects shown in the design table above, collected in 18 separate sessions, with typically 12 or 14 subjects participating in each session. Sessions took about 60 to 90 minutes (including instructions) to complete, and subjects earned about $22 on average.[6]

[6] We also varied the feedback that subjects received each period in treatments A and B. In some sessions subjects learned only the outcome of the game they played, while in other sessions they learned the outcomes of all games in the session and were provided with a continuously updated total of the cumulative frequency of outcomes experienced by all pairs in their session. We found no evidence that play was sensitive to the two types of outcome feedback, however, and therefore we pool the results for these two forms of feedback in the analysis presented in the next section.

4 Results

4.1 Human Opponents With and Without Recommendations

Figure 2 presents the overall frequency of the cooperative Down/Right choices made in each ten-period block of periods for the inexperienced sessions in which subjects played against other human opponents (treatments A and B). Although the average frequency in both treatments begins near the correlated equilibrium target of 0.6, the Down/Right choice rate declines without recommendations until it reaches the mixed strategy equilibrium of 0.4 by the end of the sessions. This average, of course, obscures some substantial differences across individual subjects, since some nearly always chose Down or Right while others nearly always chose Up or Left. What matters for the population game with random rematching is the overall expected frequency in the population. When we consider the cross-sectional rate of Down/Right choices calculated individually for each subject across the final 30 periods of the session, the mean rate of Down/Right is 0.43, which is not significantly different from the mixed strategy equilibrium of 0.4 (t-statistic = 0.78).[7]

The average frequency of the cooperative Down/Right choice also declines in the treatment with recommendations, although it declines more slowly and does not fall below 0.5 until the final 10 periods of the session. During the final 30 periods, calculated across all 80 subjects in this treatment, the mean rate of Down/Right is 0.49. This frequency is significantly different from the mixed strategy equilibrium of 0.4 (t-statistic = 3.67) and is also significantly different from the correlated equilibrium target of 0.6 (t-statistic = 4.09). A comparison of the two treatments indicates that recommendations significantly increased

[7] Throughout the results section we report cross-sectional t-statistics based on late (periods 31-60) averages calculated for individual subjects and then compared to a theoretical benchmark or across treatments. Virtually identical conclusions obtain if we employ standard errors that are calculated to be robust to unmodeled correlation of choices within clusters defined by the individual sessions.

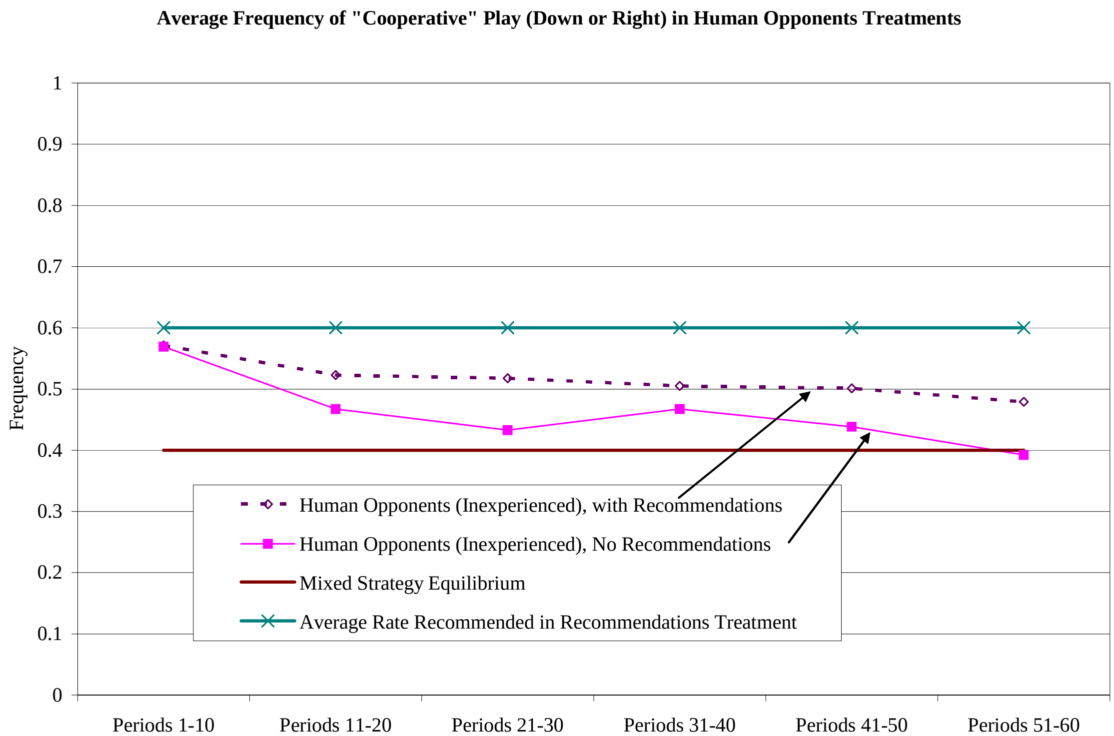

Figure 2:

the rate of cooperation (Down/Right), but at a marginal 10 percent significance level (one-tailed), with t-statistic $= 1.33$.

The average profits earned by subjects for these two treatments provides statistical evidence that parallels the choice frequency data. During the final 30 periods the cross-sectional average profit earned by subjects in the treatment without recommendations was 22.43, which is marginally significantly different (at the 8 percent level) from the mixed strategy equilibrium level of 21 (t-statistic $= 1.78$). Average profits were 25.43 during these late periods in the treatment with recommendations, which is highly significantly different from both the mixed strategy equilibrium of 21 (t-statistic $= 8.00$) and from the correlated equilibrium target of 30.6 (t-statistic $= 9.34$). A comparison across treatments concludes that recommendations significantly increased profits compared to the treatment without recommendations (t-statistic $= 3.18$). The greater statistical significance across treatments for profits, compared to the frequency of Down/Right choices, indicates that recommendations did increase coordination and were partially successful in reducing the frequency of the worst {Up, Left} outcome. For example, the {Up, Left} outcome occurs in 31.0 percent of the final 30 rounds without recommendations, compared to 19.7 percent of the final 30 rounds with recommendations. Although both are below the mixed strategy equilibrium prediction of $0.6^2 = 0.36$, recommendations clearly failed to reduce the rate of this outcome to zero.

This result obviously requires subjects to frequently disregard their recommendations. As documented more in the next subsection, during the final 30 periods the subjects in this human opponent treatment followed their recommendations only 80.4 percent of the time. The rate was substantially higher for the Up/Left recommendation (88.3 percent) than for the Down/Right recommendation (75.1 percent). This finding led us to introduce the treatments with robot opponents.

4.2 Robot Opponents

Table 2 displays the average profits earned by subjects in all treatments for the second half of the sessions (periods 31-60). Recall that in the treatments with robot opponents the subjects were fully aware that they were interacting with robots, and furthermore the instructions emphasized that the robots would always follow their recommendations. We present a separate column to indicate the average profit that subjects would earn if they always followed recommendations. The mean "recommended profit" is different in the robot opponent treatments (26.73) from the target correlated equilibrium of 30.6 because we used the same random sequence of recommendations for all subjects facing robots to reduce across-subject variation. (This explains the zero standard error for this recommended mean.) The recommended profits are lower when subjects faced robots because the random sequence of recommendations used in the robot-opponent sessions happened to have fewer Right recommendations for the robot than expected during these late periods. Therefore, our discussion and statistical analysis will focus on the ratio of actual to recommended profits (displayed in the rightmost column of the table) to normalize profits by their recommended level.

Table 2: Average Profits Across Subjects within Treatments (Periods 31-60)

Index	Treatment	Mean Profit (S.E. Mean)	Mean Profit Recomm. (S.E. Mean)	Mean Ratio Actual/Reco. (S.E. Mean)
A	Human opponents, no recommendations	22.43 (0.80)	—	—
B	Human opponents, with recommendations	25.43 (0.55)	30.67 (0.30)	0.829 (0.016)
C	Robot opponents, with recommendations, no one receives robot earnings	24.14 (0.09)	26.73 (0)	0.903 (0.003)
D	Robot opponents, with recommendations, humans receive robot earnings	24.10 (0.10)	26.73 (0)	0.901 (0.004)
E	Human opponents, with recommendations (experienced vs. humans)	24.65 (1.05)	31.04 (0.60)	0.796 (0.030)
F	Human opponents, with recommendations (experienced vs. robots)	25.78 (0.83)	30.28 (0.44)	0.855 (0.029)

The table indicates that profits relative to the recommended target increase when we control subjects' beliefs about the opponents' rationality and propensity to follow recommendations using the robot opponents. The inexperienced subjects facing human opponents (treatment index B) obtained less than 83 percent of the recommendations benchmark. By contrast, the inexperienced subjects facing robot opponents (treatment indexes C and D) obtained over 90 percent of the recommendations benchmark. This difference is statistically significant (t-statistic $= 3.87$). We conducted this test based on data pooled across the two robot opponents treatments, because the data fail to reject the hypothesis that any behavior or outcomes differ across treatments C and D (t-statistic $= 0.29$ for profits, t-statistic $= 0.52$ for rates of cooperation, and t-statistic $= 0.52$ for average rates that recommendations were followed).

The improvement in relative profit performance when subjects played against robots instead of humans can be attributed to their greater willingness to follow recommendations when playing against robots. Figure 3 displays the rates that players followed each type of recommendation across the different treatments. Subjects in all treatments followed Up/Left recommendations more frequently than Down/Right recommendations, which is analyzed later in subsection 4.4. Note that the rate subjects followed Down/Right recommendations when facing human opponents declined steadily over time. Controlling for the content of the recommendation, the figure clearly illustrates that subjects were more willing to follow recommendations on average when facing robot opponents. This difference is statistically significant (t-statistic $= 2.56$).

Of course, the treatment averages in these figures obscure some considerable variation across individuals, and some individuals in all treatments followed recommendations consistently. To highlight this heterogeneity, Figure 4 displays the cumulative distribution of the rates that individual subjects in each treatment followed their recommendations. Note that 40 to 50 percent of the individual subjects in the robot opponents treatments always followed their recommendations. Pooling across the two robot treatments, which is justified because the two treatments are not significantly different,[8] three-quarters of the subjects

[8] The Kolmogorov-Smirnov two-sample test p-value $= 0.641$. Note that this test can be applied because the individual subjects' choices in the robot-opponents data are statistically independent.

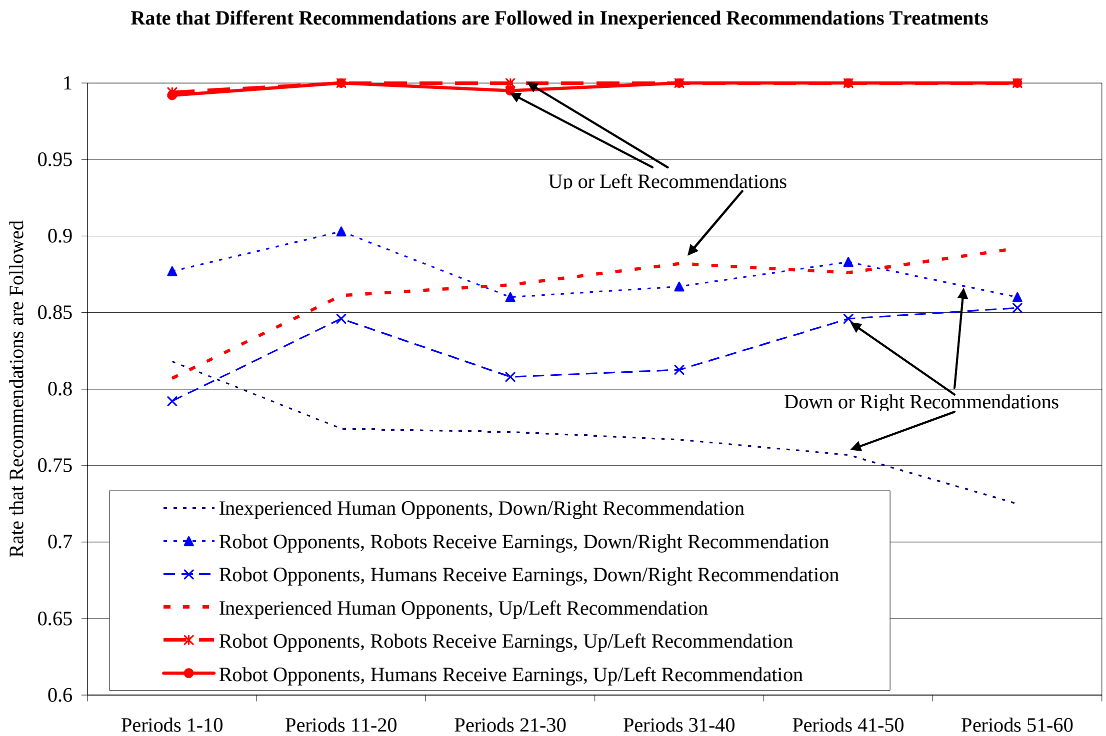

Figure 3:

playing against robots followed the recommendations at least 90 percent of the time. By contrast, less than one-quarter of the individual subjects in the human opponents treatment always followed their recommendation, and less than half followed their recommendations at least 90 percent of the time. Remember that it is a strict best response for subjects to follow their recommendation *if they believe* that their counterpart will also follow her recommendations. This figure suggests that such beliefs may not be accurate for the human opponents treatment, and subsection 4.4 analyzes reported beliefs in detail.

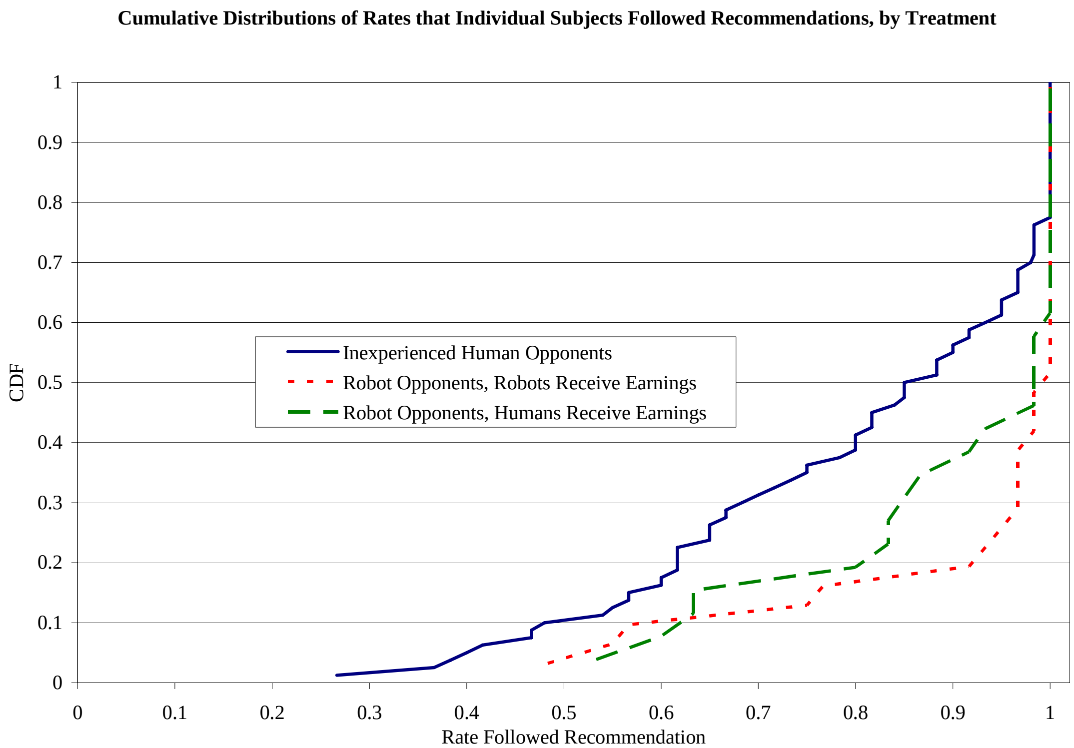

Figure 4:

4.3 Experienced Human Opponents

Table 2 above also reports the overall performance for subjects who were experienced in one of the other treatments. Regardless of whether these subjects obtained experience playing against other humans (treatment index E) or against robots (treatment index F), they still earn only about 80 to 85 percent of the recommended profit. Although profits are slightly higher when subjects received experience playing against robots, this difference is not statistically significant (t-statistic $= 1.40$). The two experienced treatments are also not significantly different in their rates of cooperation (t-statistic $= 0.54$) or in their rates that recommendations are followed (t-statistic $= 0.07$). We therefore pool these experienced treatment data in the remainder of the analysis.

The mean ratio of actual/recommended profits across the 48 experienced subjects pooled is 0.825, virtually identical to the inexperienced subjects facing human opponents (treatment index B) mean of 0.829. Thus we conclude that experience does not improve profit performance (t-statistic = 0.14). During the late periods 31-60 of both types of sessions, experienced subjects actually chose the cooperative Down/Right action slightly less than did inexperienced subjects (47.4 percent compared to 49.5 percent of the time), although this difference is not significant (t-statistic = 0.50). During these late periods the experienced subjects tended to follow recommendations more frequently than did inexperienced subjects (85.2 percent compared to 80.4 percent of the time), but again this difference is not statistically significant (t-statistic = 1.24). Overall, we conclude that when subjects play against other humans, increased subject experience does not improve the ability to implement reliably the correlated equilibrium.

4.4 Beliefs

Recommendations provide information for subjects to update their beliefs regarding their opponent's action. Specifically, if subjects believe that opponents always follow recommendations, after receiving a recommendation to play Up or Left they should report a belief that their opponent will play Right or Down with certainty because the correlated equilibrium being implemented never selects the bad {Up, Left} matrix cell; and after receiving a recommendation to play Down or Right they should report a belief that their opponent will play Right or Down with probability one-third. Recall also that we elicited beliefs from subjects and paid them for their accuracy using a proper scoring rule.

Figure 5 summarizes the beliefs data for the last 30 periods of the sessions and indicates that the recommendations did shift average reported beliefs in the predicted direction;

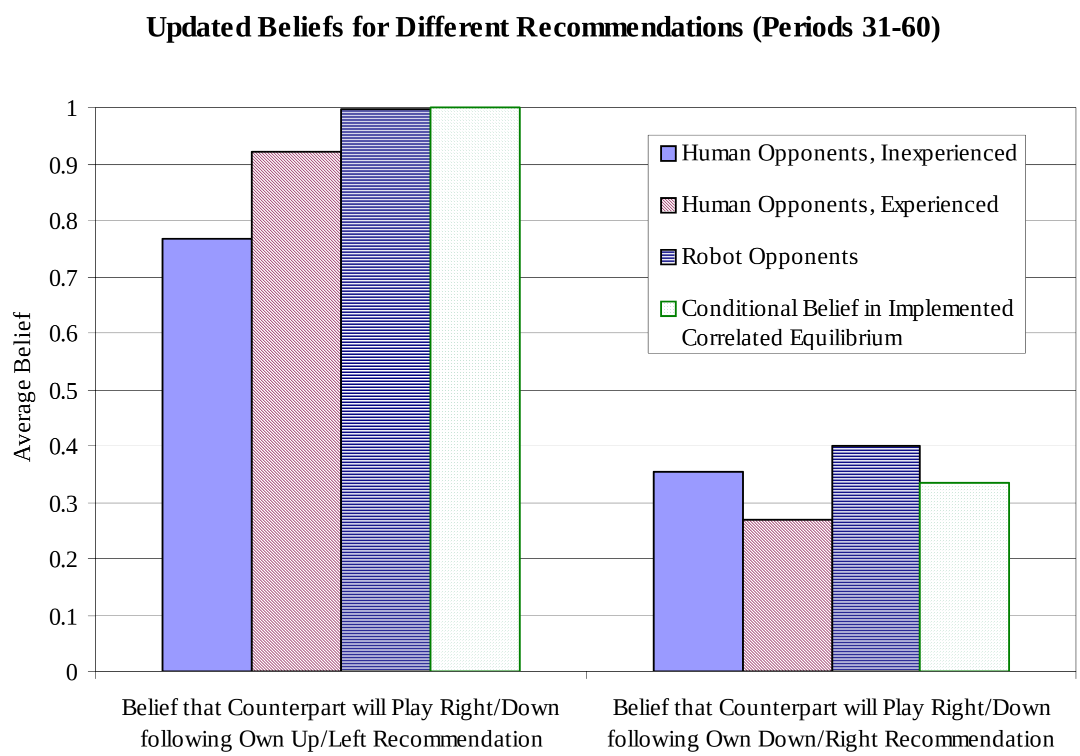

Figure 5:

however, they also indicate that subjects did not always expect a human counterpart to follow her recommendation. This is illustrated most clearly on the left side of the figure. Instead of adjusting conditional beliefs to 1 following recommendations to play Up or Left, when inexperienced subjects played against human opponents their Right/Down belief only increased on average to about 0.77. Experienced subjects report beliefs following recommendations that slightly exceed 0.9 on average, but beliefs do not shift to the prediction of 100 Right/Down except when subjects played against robots. The right side of this figure shows that average beliefs following a recommendation to play Down or Right roughly correspond to the prediction of one-third for all treatments, varying between 0.25 and 0.4.

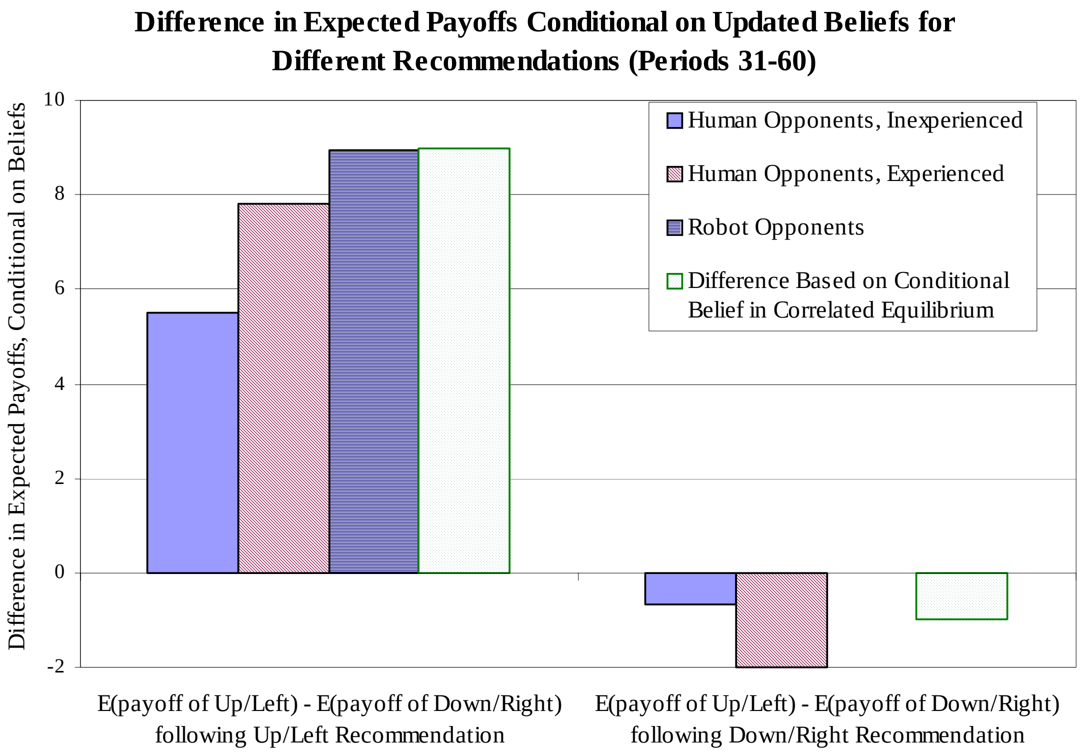

Figure 6:

This shift of beliefs following an Up/Left recommendation to less than 100 percent Right/Down when subjects played against humans is consistent with the actual rates that human subjects followed their recommendations. As shown above in Figure 3, inexperienced players facing human opponents followed Down/Right recommendations roughly 75 to 80 percent of the time; moreover, they followed this recommendation at a declining rate across the 60 periods. Experienced players facing human opponents (not shown) followed Down/Right recommendations 83 percent of the time during periods 1-20, but only 72 to 78 percent of the time during periods 41-60. Therefore the reported beliefs were reasonably accurate given the actual rates that human opponents followed recommendations.

Figure 6 summarizes the incentives that subjects have to follow recommendations given their reported beliefs. By design, incentives are much stronger to follow an Up/Left recommendation, which is displayed on the left side of this figure. The expected payoff difference for following an Up/Left recommendation in this correlated equilibrium is nine times higher than for following a Down/Right recommendation. Nevertheless, note that because beliefs do not fully adjust to 100 percent Right/Down following this recommendation, the expected payoff differential from following an Up/Left recommendation for inexperienced subjects is about one-half of the differential based on beliefs that their counterpart will always follow her recommendation. The right side of the figure indicates that subjects have considerably weaker incentives to follow a Down/Right recommendation. If subjects are more likely to make errors that have a lower expected cost, as in the quantal response equilibrium (McKelvey and Palfrey, 1995), then we would expect that Down/Right recommendations would be followed at a lower rate than Up/Left recommendations. This is exactly the pattern of behavior that our data exhibits (Figure 3).[9]

5 Conclusion

We started by asking whether players can be induced to choose *some particular* equilibrium of a game with signals through recommended play. Our results indicate that recommendations are not sufficient to induce players to switch to *any* chosen equilibrium. Subjects do not switch because they believe that their opponents will not switch. This much is clear. But

[9] If we exclude the cases of indifference or "near-indifference," defined as situations in which reported beliefs indicate that the two choices' expected profits differ by less than five percent, we find that beliefs support following Up/Left recommendations much more frequently than Down/Right recommendations. In particular, reported beliefs indicate that it is optimal to follow 90 percent of the Up/Left recommendations, but only 51 percent of the Down/Right recommendations. Both recommendations are usually followed (82 to 92 percent of the time) when beliefs indicate that it is optimal to follow them, and both recommendations are followed about 70 percent of the time even when the beliefs indicate that rejecting the recommendation is optimal.

why is this so? Given the lack of theory in this domain and given our process of data collection, we are unable to provide a complete answer to this question. Nevertheless we take a bold stride and offer some conjectures.

The reader may conclude from our data that recommendations were not totally rejected. Nearly one-quarter of subjects who played against human opponents always followed recommendations. Furthermore, our results on beliefs indicate that experienced subjects assigned more probability to opponents following recommendations than did inexperienced subjects. Moreover, these beliefs were fairly compatible with actual play of opponents. Thus we are tempted to say that, given sufficient time subjects would follow recommendations. But why don't subjects follow recommendations immediately? We provide the following conjecture. Perhaps subjects have belief hierarchies where they believe that their opponents follow recommendations but make mistakes and choose a different action with a certain probability, they believe that their opponent believes that the subject in question makes mistakes with some probability and so on. Given such a hierarchy, under certain conditions subjects are able to form a belief over their opponent's play (Brandenberger and Dekel, 1993). Given these beliefs, subjects choose their optimal actions. Based on these observed actions subjects form a different hierarchy of beliefs with potentially different probabilities of making mistakes. This different belief hierarchy determines present choice of actions. Our results seem to suggest that hierarchies put less weight on mistakes over time. From a theoretical perspective it would be interesting to know under what conditions would probabilities on mistakes (in belief hierarchies) converge to zero.

Our conjecture on mistakes, as stated above, stems from our observations on data from the robot treatments. Here, subjects almost always followed the Up or Left recommendations, whereas only around ninety percent of the subjects followed the Down or Right recommendation (see Figure 3). Note that after the Up or Left recommendation, updated probabilities imply that the robot would play Right or Down with probability one. Since the instructions clearly state that the outcome (Up, Left) will never be recommended, even a visual verification of the matrix suggests the correct updating. However, for a Right or Down recommendation subjects have to actually 'calculate' the updated probabilities. It is quite plausible for subjects to make mistakes in these calculations. Suppose this is true and assume that subjects know it. Now consider the treatments with human subjects and note that a little more than ten percent of the subjects did not follow even the Up or Left recommendation. But this is now understandable. A subject given a Up recommendation (say) knows for certain that her opponent is given the Right recommendation. But, opponents make mistakes when they are given the Right recommendation. So, it may not always be optimal for our subject to follow the Up recommendation. Similarly, for the Right or Down recommendation.

Finally, we would like to emphasize that our objective was to implement a correlated equilibrium with payoffs outside the convex hull of Nash equilibrium payoffs, and not any arbitrary correlated equilibrium. It may well be that it is easier to make subjects shift to another correlated equilibrium, such as one in which the Down/Left (9, 48) and Up/Right (48, 9) cells are recommended with equal probability.[10] But even then we would be bogged down by questions similar to the ones raised above. Why do we expect subjects to believe that their opponents will follow recommendations? Similarly, our results may also suggest that there is a set of implementable "behavioral correlated equilibria" smaller than the set of correlated equilibrium. The boundaries of this set are defined by incentives large enough

[10]The expected payoff incentive to follow recommendations in this correlated equilibrium ranges between 6 and 9, depending on the recommendation, which is similar to the incentive (9) that led Up/Left recommendations to be followed by subjects in the current experiment.

relative to "behavioral noise" arising from human tendencies to make mistakes. Determining the set of "behavioral correlated equilibria" through experiments is a desirable exercise. But then theory still has to deal with the question as to what set of correlated equilibria can be implemented through simple recommendations and what set of correlated equilibria cannot. Our paper suggests that the latter set exists.

References

[1] Aumann, R. J.: Common priors: A reply to Gul. Econometrica 66, 929-938 (1998).

[2] Aumann, R. J., A. Brandenburger: Epistemic conditions for Nash equilibrium. Econometrica 63, 1161-1180 (1995).

[3] Bolton, G., Ockenfels, A.: ERC: A theory of equity, reciprocity and competition. American Economic Review 90, 166-193 (2000).

[4] Brandenburger, A., Dekel, E. : Hierarchies of beliefs and common knowledge. Journal of Economic Theory 59, 189-198 (1993).

[5] Brandts, J., MacLeod, W. B.: Equilibrium selection in experimental games with recommended play. Games and Economic Behavior 11, 36-63 (1995).

[6] Fehr, E., Schmidt, K.: A theory of fairness, competition and cooperation. Quarterly Journal of Economics 114, 817-868 (1999).

[7] Gul, F.: A comment on Aumann's Bayesian view. Econometrica 66, 923-927 (1998).

[8] Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. Econometrica 68, 1127-1150 (2000).

[9] Holmstrom, B.: Moral hazard in teams. Bell Journal of Economics 13. 324-340 (1982).

[10] Houser, D., Kurzban, R.: Revisiting kindness and confusion in public goods experiments. American Economic Review 92, 1062-1069 (2002).

[11] Marimon, R., Spear, S., Sunder, S.: Expectationally-driven market volatility: An Experimental Study. Journal of Economic Theory 61, 74-103 (1993).

[12] McKelvey, R., Palfrey, T.: Quantal response equilibria for normal-form games. Games and Economic Behavior 10, 6-38 (1995).

[13] Mookherjee, D.: Optimal incentive schemes with many agents. Review of Economic Studies 51, 433-446 (1984).

[14] Moreno, D., Wooders, J.: An experimental study of communication and coordination in noncooperative games. Games and Economic Behavior 24, 47-76 (1998).

[15] Nyarko, Y., Schotter, A.: An experimental study of belief learning using elicited beliefs. Econometrica 70, 971-1005 (2002).

[16] Offerman, T., Sonnemans, J., Schram, A.: Value orientations, expectations and voluntary contributions in public goods. Economic Journal 106, 817-845 (1996).

[17] Ruström, E., Wilcox, N.: Learning and belief elicitation: observer effects. Working paper, University of Houston (2004).

[18] Seely, B., Van Huyck, J., Battalio, R.: Credible assignments can improve efficiency in laboratory public goods games. Journal of Public Economics 89, 1437-1455 (2005).

[19] Sharma, T., Torres, R.: Coordination in teams. Working paper, ITAM (2004).

[20] Van Huyck, J., Gillette, A., Battalio, R.: Credible assignments in coordination games. Games and Economic Behavior 4, 606-626 (1992).