

Coordination through Reputations: A Laboratory Experiment*

Donald J. Dale[†] John Morgan[‡] Robert W. Rosenthal[§]

July 1999

Abstract

One of the functions of reputations in societies is to enable coordination, and some of the most familiar instances of this occur when reputations are based on past actions. For example, a history of aggressive play is often interpreted as predicting more aggressive play in the future. This permits a pair of individuals with unequal reputations for aggressiveness to avoid wasteful conflict that would result from aggressive play by both.

This paper describes a laboratory experiment designed to explore whether, in a population of subjects, such coordination through reputation formation will emerge spontaneously, what form it will take, and how reliably it will be observed by the individuals in the population. Reputation-based coordination of the sort illustrated above did arise in all experimental sessions, but the rate of adherence to the social rule implicitly adopted by the group varied across the experimental populations as a function of treatment variables.

*We are grateful to participants at seminars at M.I.T. and Boston University for helpful suggestions, to the National Science Foundation for financial support, and to Sumon Majumdar for research assistance. Data from these experiments is downloadable (in Stata 6.0 format) at www.wws.princeton.edu/~rjmorgan.

[†]Department of Economics, Princeton University

[‡]Department of Economics and the Woodrow Wilson School, Princeton University

[§]Department of Economics, Boston University

1 Introduction

The idea that an individual might wish to condition her actions on the reputations of the people with whom she is interacting is ubiquitous in human societies, and the functions such reputational systems can serve are numerous. In this paper, we study one of the familiar functions of reputations—that of a coordination device. The basic idea is that if two individuals must make simultaneous decisions in a symmetric situation and if there is some advantage to at least one of the individuals if they act differently from each other, then any external cues, if agreed upon either tacitly or explicitly, can serve to direct them toward a coordinated asymmetric outcome. Reputations can provide such external cues. As an analogy, many species of animals are observed to respect “pecking orders” which peacefully desymmetrize situations that might otherwise result in conflict. (A typical example is males competing for the opportunity to copulate with all the fertile females in a herd. The top male often wins without a fight). For a human example, an individual (or firm) having the reputation of a bully may be able to avoid conflict and prosper because potential competitors either defer to the bully when in situations of conflict or avoid such situations to begin with.

There are complications even in these relatively simple examples, however. Is the top animal in the pecking order the individual with the toughest appearance (e.g., largest size, sharpest antlers, brightest color) or the one with the best record in past conflicts? How many instances of bullylike behavior does it take for the reputation of bully to be acquired and how long does it last? In short, what is the nature of the external cue? In real-world situations the natures of external cues can differ for reasons having much more to do with example-specific details than with the coordinating role of the cue. For example, the ability to remember the identity of the strongest or weakest animal in one’s herd might have survival value in the face of threats from external predators or might be useful in finding food. Similarly, bully reputations might correlate with aggressive tendencies toward external foes; and it might therefore pay either to align with or against bullies in one’s own society.

To learn about how humans use reputations as coordinating devices, therefore, it seems useful to explore the question in the laboratory, where many of these other factors can be controlled for, at least to some extent. In this paper we report on one such experiment. Our hope is that the aspects of everyday behavior that concern the coordinating role of reputations are

maintained in experimental behavior despite the abstract setting of the laboratory environment. We know of no previous experiments on this subject.

We report here on the experimental results from a variety of sequenced 2×2 coordination games with successively rematched player pairs, where prior to play in each round each pair of subjects was presented with a pair of summary signals, called “current positions,” that reveal information about the recent play of both members of the pair. We did not describe the positions to the subjects as reputations and did not suggest to the subjects that conditioning on positions was appropriate. We merely gave the group of subjects the opportunity to develop, collectively but implicitly (since they were given no opportunities to communicate directly with each other), a coordination mechanism by providing to each matched pair of players the common knowledge of their respective current positions.¹ Our interest was in whether, how, and how reliably, our experimental societies would use the position pairs as coordinating devices and in which of the many possible position-based coordinating rules they would adopt.

More specifically, the subjects in a session played repeatedly the same variant of either

$$\begin{array}{|c|c|} \hline 0, 0 & x - y, x + y \\ \hline x + y, x - y & 0, 0 \\ \hline \end{array} \text{ or } \begin{array}{|c|c|} \hline x, x & x - y, x + y \\ \hline x + y, x - y & 0, 0 \\ \hline \end{array},$$

where $x > y \geq 0$. Identifying the top row with the left column, all such one-shot games are seen to be symmetric, but two of their Nash equilibria – the ones that use pure strategies – are not symmetric. Furthermore, though the third Nash equilibrium in each such game is symmetric, it involves an inefficient mixed-strategy combination. Each subject’s position is a signal about how frequently in the recent past she has played the bottom row (equivalently, the right column) which we might think of as the more aggressive action.

Our prior guess was that the incentives to coordinate on the pure Nash equilibria of the one-shot games would be sufficiently strong that the subjects would find ways to achieve them quite frequently. Furthermore, we suspected that the organizing rule with the most salience would be that in cases where a matched pair finds itself in unequal positions the player having the “tougher” current position (i.e., the position signalling more frequent recent play of

¹Of course there might be an experimenter effect: subjects might try to use their positions to coordinate solely because they are prominently available in the experiment.

the aggressive action) would play the aggressive strategy while the player with the “softer” current position would choose the other action. This seems to accord with everyday intuition: individuals who have reputations formed from their records of previous actions are commonly expected to continue as in the past. But why this rule is so intuitive is itself not obvious; after all, the reverse rule (tougher reputation types playing soft and softer types playing tough) coordinates similarly. In addition, there are many other patterns which, though more complicated in that the identity of the player chosen to play tough depends on more details of the state pairs, would also achieve coordination. Perhaps it is the simplicity of the intuitive rule that makes it salient (although the reverse rule seems equally simple). Perhaps it is the fact that people who play according to the simplest of all rules – constantly the same pure strategy – end up in the extreme positions and therefore are not observed to violate the intuitive rule. Perhaps the intuitive rule has a kind of Darwinian survival value, in which case the reason it seems intuitive to us is that it is the rule used in the societies we know about; namely, the ones that have survived.² And perhaps our prior guesses would be proved wrong – what seems focal to us is not necessarily focal to everyone, and if there is heterogeneity in intuition across the population, this may make adoption of any one rule difficult for the subject population.³

Our experimental findings have turned out to be quite interesting. We ran sixteen sessions, all with initially inexperienced subjects, for 100 rounds each. Ten of our sessions were based on parametric variants of “Battle of the Sexes,” or BOS, the first game form above; and six used parametric variants of “Hawk-Dove” or “Chicken,” or HD, the second game form above. Examining the last 50 rounds of each of the sessions, we find that, to the extent that any one simple rule explains play, it is overwhelmingly the intuitive rule. In the BOS sessions, when a player is matched with a co-player in an unequal position, she plays according to the intuitive rule 90.6% of the time. The intuitive rule is somewhat less successful in the HD sessions, but it still predicts 78.9% of choices correctly for players matched with an unequally positioned opponent. Both satisficing and altruistic motives could be explanations for this difference across game forms: The less aggressive strategy is relatively safe in HD, and if its payoff against either strategy of the opponent is good

²This hypothesis is explored theoretically in Rosenthal and Landau (1979). We make use of their modeling framework here.

³There is an extensive literature on focal points. The classic treatment is Schelling (1960).

enough, why risk it for a small gain? This safe strategy is also kindly toward the opponent no matter what he does.

The compliance percentages from unequal positions listed above are averages over treatments for each game form in which the y -payoff parameter was varied. In both BOS and HD treatments, increasing the y -payoff parameter unambiguously reduced the predictive power of the intuitive rule. When $y = 0$ in BOS, the game becomes a pure coordination game and the intuitive rule correctly predicted 97.8% of choices; this declined to 90.8% when a small positive y was used and still farther to 81.4% in the treatment having the highest y value. The situation was similar for HD: With the same small y value as in BOS, the intuitive rule correctly predicted 91.0% of choices. With the larger y value, the intuitive rule predicted 72.9% of HD choices.⁴

In fact, the majority of departures from the intuitive rule, especially in the BOS sessions, were by players who deviated to the aggressive strategy. This suggests a calculation involving a sacrifice of current payoff for improved future payoffs. In order to quantify such trade-offs and thereby attempt to explain the varying predictive power of the intuitive rule across experimental sessions, however, it is useful to view the experiment through the lens of a well-specified mathematical model.

The model of Rosenthal and Landau (1979) is a relatively simple lens through which to view this experiment. It depicts a society consisting of a large number of individuals who are repeatedly randomly matched to play a game (having von-Neumann-Morgenstern utility payoffs) over an infinite time sequence, where all individuals are assumed to discount according to a common fixed factor.⁵ Actions produce effects on payoff-irrelevant individual state variables, and matched pairs of individuals know their own and opponents' individual states, just as in our experiment the matched pairs of subjects know both positions. A steady state equilibrium of this model is a stationary Markov rule that assigns to each ordered state pair ((own, opponent's)) a strategy (possibly mixed) in the stage game such that each used action is a best response in each state pair, taking the future into account as follows: The states occupied by an individual's future matched opponents will be i.i.d. draws from the unique invariant distribution over individual states generated by the Markov rule, and all individuals will play according

⁴HD with $y = 0$ is a game with different equilibrium structure, so we did not try any such treatments.

⁵Rosenthal and Landau restricted their analysis to a particular parametrization of HD, but their model is usable more generally.

to the rule in the future. Furthermore, an extension of the intuitive rule (so that it also makes choices in equal-position situations) generates one of the steady state equilibria of the model when the common discount factor is not too high relative to the parameter y ; and the smaller is y , the larger is the set of common discount factors for which the intuitive rule generates a steady state equilibrium.

The applicability of this model to the experiment is evidently suspect because of its simplifying assumptions of an infinite future, discounting, money payoffs being von-Neumann-Morgenstern utilities, and an infinite player set. Despite these, the equilibrium of the model corresponding to the intuitive rule seems to organize the data reasonably well. (We shall discuss later why the assumptions above might not seriously misrepresent our subjects' behavior.) So, although the relationship between the model and the experiment seems crude, it fits the data reasonably well in many respects, provides at least a hint of an explanation for the observed effect of the y parameter, and therefore leads us to the following tentative conclusions.

1. Of the model's many steady state equilibria, only the one corresponding to the intuitive rule approximates play observed in the laboratory.
2. Experimental play of the intuitive rule is most reliable in BOS with small y parameter and least reliable in HD with high y parameter; and, within each game form, reliability is monotone in y .
3. Since for any fixed common discount factor and fixed game form the intuitive rule generates a steady state equilibrium only when y is sufficiently small, the relatively high noise levels in the data for the experimental sessions with larger y values reflect the increased incentives that relatively patient subjects have to deviate from the prescriptions of the intuitive rule.⁶

We hasten to add that our claim that the model fits the experimental data reasonably well is inherently subjective. The intuitive rule, as formulated here, allows for no variance in play when players in unequal positions are matched. Since there is some variance in the experimental data, statistical tests would reject the hypothesis that the model could have generated the data. Furthermore, we have run only 16 sessions; and these encompass a

⁶We will discuss later why we think degrees of impatience might matter in our experiment even though all payments are made at the end of the experimental session.

variety of treatments. In addition, the model predicts subject play least well in situations where a subject and her matched opponent are in equal positions; i.e., in the situations where the intuition for the intuitive rules seems murkiest. Thus our findings could be interpreted more as supportive of the intuitive rule than of the intuitive rule together with the model in which we have chosen to embed it.

Because many features of the experiment and many of the questions that can potentially be addressed using its data are complex, we intend to present our findings in stages. This paper, the first such stage, is devoted to the “long run,” or what happens after play settles down. We believe that our subjects have learned sufficiently about their environment by round 50 and that their behavior after that is in the aggregate approximately stationary. (We will present evidence for this as well as for the absence of end-game effects in the data.) Accordingly, in this paper we will restrict our attention almost entirely to the data from rounds 51-100 of each session. We will then compare the experimental results for these rounds with certain steady state equilibria of the theoretical model, leaving for future analysis such interesting questions as why the intuitive rule is apparently focal and how the populations of subjects learn to play it.


The rest of this paper is organized as follows. Section 2 describes the experimental design and procedures. Section 3 presents the theoretical model to be adapted to the experiment. Section 4 presents the results of our experiment and compares them to the theoretical predictions. Section 5 concludes.

2 Experimental Design and Procedures


The experiment consisted of 16 sessions conducted at Princeton University in Spring 1999. Twelve volunteer subjects (Princeton undergraduates and/or Masters students from the Woodrow Wilson School) participated in each session, with no subjects taking part in more than one session. Subjects were recruited by e-mail invitations to participate in a decision making experiment. They were promised between \$5 and \$35 for a session lasting approximately 60 minutes. In fact, no session lasted longer than 60 minutes; the lowest payoff actually earned in the experiment was \$12.70; and the highest amount earned was \$27.25. Even the lowest amount is considerably more than the hourly outside earning rate of the typical student. The subjects were paid their winnings individually in cash at the end of the session.

The following procedures were common to all sessions. At the beginning of a session the subjects were seated at computer terminals in a single room and given a set of instructions. These instructions were read aloud, and the subjects were then given an opportunity to ask questions. (See the appendix for a sample copy of the instructions used.) Subjects then played the same symmetric 2×2 normal-form game 100 times in succession. No communication between subjects was permitted, and once play began all choices and information were transmitted via the computer terminals. In each round, each subject was randomly and anonymously paired with another individual in the session, with equal probability given to all eleven others in each round. Subjects chose between the two possible actions, one labeled “ \rightarrow ” and the other labeled “ \leftarrow ”.


Prior to choosing an action in each round, each subject was told her own current position, the current position of the person with whom she was currently matched, and the number of subjects currently at each position. There were five possible current positions given by tick marks on a horizontal line:



At the start of the first round, all subjects began in the middle position. That is, each subject’s initial current position was



For round 2, if a subject had previously chosen “ \rightarrow ,” then her current position would be shifted one tick mark to the right. Likewise, if in round 1 she had chosen “ \leftarrow ,” her current position would be shifted one tick to the left. In all subsequent rounds, a similar application of this scheme determined all subjects’ current positions with two exceptions. If a subject is at the rightmost tick mark and chose the action “ \rightarrow ,” then she stayed at the rightmost tick mark. Similarly, if a subject was at the leftmost tick mark and chooses action “ \leftarrow ,” then she remained at the leftmost tick mark. Thus, a player’s current position provided a signal about recent play. Figure 1 displays a typical first-round screen including current positions. For purposes of exposition, we shall number the five positions as follows:



but no such delineation was made for the subjects participating in the experiment.

Following the selection of an action by all subjects in a given round, the payoff cell jointly selected by a subject and the matched opponent was illuminated on the subject’s screen. Subjects then had to click their computer mice to continue to the next round. Subjects were kept continuously informed of their cumulative earnings throughout the session. At its conclusion, subjects were asked to fill out a post-experiment questionnaire explaining how they played during the experiment and were paid their winnings privately.

Early rounds of the experiment tended to be completed relatively slowly with more than a minute sometimes elapsing between the completion of successive rounds.⁷ By round 20, however, the speed with which decisions were being made increased substantially. The time elapsing between rounds was by then typically not more than about 15 seconds. This pace remained relatively constant for the remainder of the session; however, in some sessions play slowed down again in the last few rounds. During the early rounds, subjects appeared to be testing various strategies and generally getting a “feel” for the game. In later rounds, subjects seemed mostly to be just following whatever decision rules they had settled on. Occasionally there would be a round that was delayed by a subject seemingly deciding to re-evaluate her strategy.

Treatments

As discussed earlier, the treatments all involved a stage-game money payoff matrix of the form

z, z	$x - y, x + y$
$x + y, x - y$	$0, 0$

where $x > y \geq 0$.⁸ When $z = 0$, the game is a BOS variant, and when $z = x$ the game is an HD variant.⁹

⁷No play in a round was possible until all subjects had made choices in the preceding round.

⁸As is customary, the first number in the cell is the payoff to the row player.

⁹To be precise, in our experimental implementation we added 5 cents to both payoffs in each cell of the matrix to guarantee that all subjects received at least \$5 for their time.

For all treatments we fixed x at 20 cents, and we varied the amount y across treatments within each game form (BOS or HD). Thus, we shall use the shorthand BOS(y') to denote a session where the BOS payoff matrix was used and the y parameter was assigned the value y' (in cents); similarly for HD(y').¹⁰ The following table summarizes the sessions we ran:

Treatment Name	# of Sessions
HD(5)	2
HD(10)	4
BOS(0)	2
BOS(5)	6
BOS(10)	2

Table 1: Summary of Treatments

(We ran more sessions for HD(10) and BOS(5) because those treatments exhibited more cross-session variation.) For all of these treatments, the intuitive rule predicts the same behavior when subjects occupying different positions encounter one another. When the members of a matched pair occupy the same position, the prediction of the equilibrium of the theoretical model corresponding to the intuitive rule depends on both the assumed discount factor and the parameter y .

Other Design Considerations

Two other aspects of the experimental design are worth mentioning.

1. To keep the environment as simple as possible, we chose not to try to induce theoretically risk-neutral preferences through a binary lottery (e.g., Roth and Malouf (1979)) despite the fact that the games we study have mixed equilibria and more than two levels of payoffs.¹¹

2. To control for the possibility that the labeling of actions might affect subject play, in half the sessions within each treatment we labeled the aggressive action “ \rightarrow ” and in the other half “ \leftarrow ”.¹² To test whether these labels

¹⁰Thus Figure 1 is for HD(10).

¹¹An additional consideration in not employing this procedure is that its success in practice in inducing risk-neutral preferences has been mixed. Selten *et al.* (1995) report on experiments showing that making direct payments for actions may actually work better than binary lotteries in inducing risk-neutral behavior.

¹²In BOS(0), where the actions are equally aggressive, the statement does not apply.

had any effect on subject choices, we conducted a t-test of the proportions of strategically equivalent choices made under the two different labelings for each treatment. As an example, a player whose current position was 4 and who was matched with a 0 and chose \rightarrow under one labeling would be deemed to be acting strategically equivalently to a 0 matched with a 4 choosing \leftarrow under the reverse labeling.¹³ For a given treatment, we tested whether the proportion of right choices in the first case was equal to the proportion of left choices in the second for each strategically equivalent reputation pair. We find no significant labeling effect at the 5% level in the last 50 rounds of play for any pair in any treatment. As a consequence, in the following exposition we pool all sessions for a given treatment and rename the positions and choices in half the sessions so that “ \rightarrow ” now always refers to the aggressive choice.

3 Theory

To provide some basis for judging the predictive power of the intuitive rule with respect to the experimental results, we adapt the Rosenthal-Landau model to our experimental setting.

For the one-shot game

	\leftarrow	\rightarrow
\leftarrow	z, z	$x - y, x + y$
\rightarrow	$x + y, x - y$	$0, 0$

with 5 positions (labelled 0, ..., 4 from left to right) and common discount factor β , a **steady state equilibrium** is a triple (π, P, v) where

$\pi = (\pi_0, \pi_1, \pi_2, \pi_3, \pi_4)$ is the population’s steady state relative-frequency distribution over the positions;

$P =$

$$\begin{pmatrix} p_{00} & \dots & p_{04} \\ \cdot & \dots & \cdot \\ p_{40} & \dots & p_{44} \end{pmatrix},$$

¹³Recall that the subjects were also kept apprised of the current population frequency distribution over positions. For this test, as for all others in this paper, we pooled the data across these different frequency distributions.

where p_{ij} is the probability that each individual assigns to the play of \rightarrow when own position is i and opponent's position is j ; and

$v = (v_0, v_1, v_2, v_3, v_4)$ is the vector of equilibrium present values.

The equilibrium conditions are:

$$\pi_0 = \pi_1 \sum_j \pi_j (1 - p_{1j}) + \pi_0 \sum_j \pi_j (1 - p_{0j}) \quad (1)$$

$$\pi_i = \pi_{i+1} \sum_j \pi_j (1 - p_{i+1,j}) + \pi_{i-1} \sum_j \pi_j p_{i-1,j} \quad (2)$$

for $i = 1, 2, 3$; and

$$\pi_4 = \pi_4 \sum_j \pi_j p_{4j} + \pi_3 \sum_j \pi_j p_{3j}. \quad (3)$$

Adopting the conventions $4 + 1 \equiv 4$ and $0 - 1 \equiv 0$,

$v_i =$

$$\sum_j \pi_j [p_{ij} ((x + y) (1 - p_{ji}) + \beta v_{i+1}) + (1 - p_{ij}) ((x - y) p_{ji} + z (1 - p_{ji}) + \beta v_{i-1})]$$

$$i = 0, \dots, 4. \quad (4)$$

Finally $\forall i, j$, $p_{ij} = 1$ if

$$(x + y) (1 - p_{ji}) + \beta v_{i+1} > (x - y) p_{ji} + z (1 - p_{ji}) + \beta v_{i-1};$$

and $p_{ij} = 0$ if

$$(x + y) (1 - p_{ji}) + \beta v_{i+1} < (x - y) p_{ji} + z (1 - p_{ji}) + \beta v_{i-1};$$

Of course, we must also have

$$0 \leq p_{ij}, \pi_i \leq 1 \quad \forall i, j, \quad \text{and} \quad \sum_i \pi_i = 1.$$

(1)-(3) require that π be an invariant distribution for the Markov chain generated by the given matching process and decision rule P . (4) defines v as the vector of present values for a player finding herself in each position

in the steady state, and the strict inequalities require that best responses be played with probability one.

To illustrate, i.i.d. play of the symmetric one-shot equilibrium

$$\left(\frac{x + y - z}{2x - z}, \frac{x - y}{2x - z} \right)$$

generates a steady state equilibrium in which

$$p_{ij} = p \equiv \frac{x + y - z}{2x - z} \quad \forall i, j, \quad v_i = \frac{x^2 - y^2}{(1 - \beta)(2x - z)} \quad \forall i, \text{ and}$$

$$\pi = \left(\sum_{i=0}^4 p^i (1 - p)^{4-i} \right)^{-1} ((1 - p)^4, p(1 - p)^3, p^2(1 - p)^2, p^3(1 - p), p^4).$$

For BOS, $z = 0$ and $y < x$; so $p \geq 1/2$, and the components of the π -vector are nondecreasing. For HD, $z = x$, so $p = y/x$; and the π -distribution is nondecreasing or decreasing according to whether $y \geq x/2$ or not. The value of β plays no role in this steady state equilibrium since players foresee no future implications of their current actions.

As another illustration, the reverse intuitive rule is consistent with a steady state equilibrium for all the BOS and HD parametrizations of this experiment. Here the value of β does enter the calculations, but as far as we have checked there seems to be a reverse-intuitive steady state equilibrium for each of the treatments for any $\beta \in [0, 1)$.¹⁴ A typical π -vector for such an equilibrium is \cap -shaped. For instance, for BOS(10) the steady state equilibrium when $\beta = .99$ has

$$p_{ij} = \begin{cases} 0 & \text{if } i > j \\ 1 & \text{if } i < j \end{cases}$$

$$(p_{00}, p_{11}, p_{22}, p_{33}, p_{44}) \cong (.50, .41, .49, .54, .58) \text{ and}$$

$$(\pi_0, \pi_1, \pi_2, \pi_3, \pi_4) \cong (.26, .44, .26, .05, .00).$$

There are many other steady state equilibria for the models. Most of these, like the two above are grossly at odds with the experimental data. There are some variants of the intuitive rule, however, that generate steady state equilibria having qualitative characteristics similar to those of the equilibria that are generated by the intuitive rule. More on this later.

For the rest of this section, we describe steady state equilibria corresponding to the intuitive rule.

¹⁴This may be an artifact of the particular treatments chosen for the experiment. See Rosenthal and Landau (1979).

3.1 The BOS Model

The intuitive rule has

$$p_{ij} = \begin{cases} 1 & \text{if } i > j \\ 0 & \text{if } i < j \end{cases} .$$

To complete the specification of an equilibrium, we need to find $\{p_{ii}\}$ for which the π and v that solve (1)-(3) and (4) satisfy the best-response conditions. The tables below give the unique approximate $\{p_{ii}\}$ and π values that do this for each of the three treatments for a variety of discount factors. In each case, the highest discount factor listed is close to the maximum discount factor for which the intuitive rule forms an equilibrium.¹⁵

β	p_{00}	p_{11}	p_{22}	p_{33}	p_{44}
.4	.528	.535	.500	.465	.472
.5	.536	.547	.500	.453	.464
.6	.545	.561	.500	.439	.455
.75	.561	.585	.500	.415	.439
.85	.573	.604	.500	.396	.427
.995	.594	.638	.500	.362	.406

Table 2: Steady State Play of \rightarrow in BOS(0) from Equal Positions

β	p_{00}	p_{11}	p_{22}	p_{33}	p_{44}
.4	.677	.707	.669	.635	.621
.5	.695	.740	.693	.644	.622
.6	.715	.783	.730	.662	.625
.75	.755	.880	.829	.722	.642
.85	.794	.987	.961	.818	.675

Table 3: Steady State Play of \rightarrow in BOS(5) from Equal Positions

¹⁵The intuitive rule forms a steady state equilibrium for every nonnegative discount factor between zero and the maximum for each treatment.

β	p_{00}	p_{11}	p_{22}	p_{33}	p_{44}
.4	.826	.877	.833	.797	.765
.5	.850	.928	.879	.825	.774
.6	.879	.993	.944	.868	.789

Table 4: Steady State Play of \rightarrow in BOS(10) from Equal Positions

Notice that for those cases where the same discount factor appears in more than one of the tables above, holding that discount factor and the position i fixed and letting the y parameter increase always results in an increased value for p_{ii} .

β	π_0	π_1	π_2	π_3	π_4
.4	.348	.108	.088	.108	.348
.5	.347	.109	.088	.109	.347
.6	.346	.110	.089	.110	.346
.75	.343	.112	.091	.112	.343
.85	.341	.113	.092	.113	.341
.995	.338	.115	.095	.115	.338

Table 5: Steady State Position Frequencies in BOS(0)

β	π_0	π_1	π_2	π_3	π_4
.4	.316	.112	.086	.099	.387
.5	.313	.113	.087	.099	.388
.6	.309	.113	.087	.100	.391
.75	.300	.113	.088	.100	.400
.85	.290	.112	.087	.100	.414

Table 6: Steady State Position Frequencies in BOS(5)

β	π_0	π_1	π_2	π_3	π_4
.4	.288	.111	.080	.082	.438
.5	.283	.111	.080	.082	.443
.6	.278	.111	.080	.080	.452

Table 7: Steady State Position Frequencies in BOS(10)

Notice that in each of the three treatments, the π -distribution is U-shaped and not very sensitive to changes in the discount factor.

3.2 The HD Model

Approximate values of $\{p_{ii}\}$ and π which together with the $\{p_{ij}\}_{j \neq i}$ of the intuitive rule generate steady state for various β -values for the HD treatments are given in the tables below.

β	p_{00}	p_{11}	p_{22}	p_{33}	p_{44}
.5	.359	.435	.390	.378	.314
.6	.396	.516	.472	.442	.340
.75	.473	.708	.697	.622	.410
.8	.506	.802	.820	.725	.451
.85	.544	.915	.979	.865	.507

Table 8: Steady State Play of \rightarrow in HD(5) from Equal Positions

β	p_{00}	p_{11}	p_{22}	p_{33}	p_{44}
.5	.688	.844	.781	.744	.618
.6	.740	.970	.924	.859	.665

Table 9: Steady State Play of \rightarrow in HD(10) from Equal Positions

Again (at least for the one case of a discount factor in common) each p_{ii} value increases in y .

β	π_0	π_1	π_2	π_3	π_4
.5	.389	.095	.085	.113	.318
.6	.376	.100	.086	.114	.325
.75	.351	.101	.088	.115	.346
.8	.340	.101	.087	.115	.357
.85	.327	.100	.0860	.113	.374

Table 10: Steady State Position Frequencies in HD(5)

β	π_0	π_1	π_2	π_3	π_4
.5	.310	.110	.087	.101	.392
.6	.297	.109	.086	.098	.410

Table 11: Steady State Position Frequencies in HD(10)

Here the π - distribution is more sensitive to changes in β .

4 Results

4.1 Preliminary Analysis

In this section, we examine whether observed behavior (in the last 50 rounds) appears qualitatively consistent with the intuitive rule overall. First, we present summary statistics on the frequency of choices that accord with the intuitive rule when matched positions are unequal.

Treatment	% Compliant
BOS(0)	97.8
BOS(5)	90.8
BOS(10)	81.4
HD(5)	91.0
HD(10)	72.9
All treatments	85.9

Table 12: Compliance with Intuitive Rule

As Table 12 makes clear, in all treatments, levels of compliance with the intuitive rule are considerably in excess of what might occur by chance. Furthermore, the compliance rates within each game form are monotonically decreasing in y . This may be formalized by testing for equality in mean compliance rates across treatments. Table 13 below presents t-statistics for pairwise comparisons of compliance rates. In all cases, differences in the y parameter are significant at the 5% level.

Treatment	t-statistic
BOS(0) vs BOS(5)	6.88
BOS(5) vs BOS(10)	7.35
HD(5) vs HD(10)	10.87

Table 13: Differences in y parameter

Interestingly, if we hold fixed the value of the y parameter and compare compliance rates for different z values, we find that there is a significant difference in the compliance rates of BOS(10) as compared to HD(10) ($t = 4.59$), but no significant difference between BOS(5) and HD(5) ($t = 0.18$). Thus, compliance rates seem to be more sensitive to variations in y than to variations in z .

To see roughly how much information about current positions affects play, it is useful to compare our compliance rates to the results of Cooper, et al. (1994). They report results of the last eleven rounds of three round-robin experiments where each session consisted of eleven players playing BOS(10). Out of 165 plays, only 41% resulted in coordinated off-diagonal outcomes. Cooper, et al. do not report choice data by individuals, but their aggregate statistics are roughly consistent with the predictions of i.i.d. draws from the symmetric mixed strategy equilibrium of the game above in which each player randomizes with probability $3/4$ on the aggressive action. To see the comparison another way, in our BOS(10) treatment 58.5% of the joint choices overall and 65.6% of the joint choices from unequal positions resulted in off-diagonal outcomes. These numbers are considerably higher than 41%. This comparison suggests that information about current positions is probably playing a large role in facilitating coordination in our experiment.¹⁶

¹⁶Cooper et al. also report on a treatment in which prior to each play one of the players was permitted (but not forced) to announce the listed name of one of the strategies, with

To get a feel for the evolution of play in our experiment, Table 14 presents compliance percentages (from unequal positions) for specific rounds. Here we see that while the intuitive custom is, to some extent, learned by the subjects; nonetheless, a considerable proportion of early moves is already consistent with the intuitive rule. This suggests that the intuitive rule might indeed be intuitive or instinctive for many subjects.

Treatment	Rd. 2	Rd. 51	Rd 100
HD(5)	100	100	100
HD(10)	63.6	68.8	78.1
BOS(0)	83.3	100	100
BOS(5)	63.3	86.5	92.6
BOS(10)	88.9	94.4	100
All treatments	71.8	87.3	91.7

Table 14: Compliance Over Time

Several aspects of Table 14 are worth highlighting. First, more than 2/3 of subject choices were already consistent with the intuitive custom in Round 2 of the experiment.¹⁷ By round 51, overall compliance rates increased farther to around 87%. In two treatments, HD(5) and BOS(0), the data were perfectly consistent with the intuitive rule in this round. And, over the last 50 rounds of the game, there was only a slight upward trend in compliance rates. A feeling for the stationarity of compliance with the intuitive rule over the last 50 rounds may also be obtained from Figures 2-7, which depict average compliance rates across treatments over all rounds. As Figures 2-7 also illustrate, there are no obvious endgame effects in the compliance rates.

The theoretical model hypothesizes that individuals are playing stationary strategies in an infinitely repeated game. A necessary condition for the data to be consistent with this hypothesis is that the empirical frequency distributions of current positions be constant. To test this, we divided the

no previous discussion of what signal such announcements might convey. In 157 of the 165 observations, the designated player announced the name of his aggressive strategy; and in 154 of these the pair successfully coordinated on the favored outcome of the announcer. Evidently, the subjects had little difficulty in appreciating and responding to this particular desymmetrizing opportunity.

¹⁷Since all subjects begin in the same current position, round 1 provides no compliance data.

observations of current positions into two groups: rounds 51-75 and rounds 76-100. The theoretical prediction is that the empirical frequency of current positions should be the same in rounds 51-75 as in rounds 76-100. Table 15 below reports p-values obtained from Fisher exact tests of this null hypothesis. As the table shows, the null hypothesis is rejected at the 5% significance level in one of the HD(10) sessions (session 4) and both of the BOS(0) sessions, but not in the remaining 13 sessions. In the sessions where the null hypothesis is rejected, much of the variation appears to occur in rounds 51-60; thus, if instead we examine only the last 40 rounds divided evenly into two groups (third column of Table 15), we find that the null hypothesis is not rejected at the 5% level in any of the previously rejected sessions.¹⁸

Session	Treatment	p-values	
		Rds 51-100	Rds 61-100
1	HD(10)	.578	.184
2	HD(10)	.288	.200
3	HD(10)	.080	.507
4	HD(10)	.008*	.147
5	BOS(10)	.257	.163
6	BOS(10)	.349	.979
7	BOS(5)	.943	.976
8	BOS(5)	.127	.837
9	BOS(5)	.127	.278
10	BOS(5)	.694	.583
11	BOS(5)	.916	.382
12	BOS(5)	.692	.568
13	HD(5)	.119	.034*
14	HD(5)	.777	.446
15	BOS(0)	.049*	.437
16	BOS(0)	.010*	.279

* denotes significant at 5% level

Table 15: Fisher Exact Tests of Stationarity

To summarize, the introduction of current position information clearly increases the frequency with which coordination is achieved relative to a

¹⁸We reject at the 5% level in only one session, a different one, using this methodology.

similar experiment where this information is absent. Moreover, the form this coordinating behavior takes appears to be qualitatively consistent with the predictions of the intuitive custom and relatively stationary for the last half of each session. To see how well the data fit the more detailed predictions of the theoretical model, we shall have to look more closely at the data for each treatment.

4.2 BOS Results

We turn first to more detailed analyses of the results from the BOS treatments.

BOS(0)

In BOS(0), there is no difference in the aggressiveness of the two actions. Here, current positions merely serves as a signal of the tendency of a particular player to be choosing the action labelled \rightarrow rather than the one labelled \leftarrow . Perhaps as a consequence of this, the compliance rates were highest for this treatment. To see in more detail the patterns of compliance, it is useful to examine the proportion of plays where the \rightarrow action was chosen as a function of the current position pairs. This information is presented for the last 50 rounds of the two BOS(0) sessions in Table 16, where the entries are the percentages of instances of \rightarrow play from each position pair and the numbers in parentheses are the totals of observed instances of the respective position pairs. (Thus, for instance, the numbers in row 0 column 4 indicate that 99.4% (i.e., 172) of the 173 instances of subjects in position 4 facing an opponent in position 0 played \rightarrow .)

		Player's Own Current Position				
		0	1	2	3	4
Opponent's Current Position	0	50.1 (126)	97.6 (42)	100 (35)	100 (44)	99.4 (173)
	1	11.9 (42)	50.0 (8)	100 (10)	100 (9)	100 (43)
	2	0.0 (35)	0.0 (10)	50.0 (6)	100 (8)	100 (33)
	3	0.0 (44)	0.0 (9)	0.0 (8)	75.0 (4)	86.3 (51)
	4	0.6 (173)	0.0 (43)	0.0 (33)	9.8 (51)	64.2 (160)

Table 16: Play of \rightarrow in BOS(0)

In the case of players in unequal positions, as already noted the intuitive prediction does remarkably well. Indeed, out of 896 instances of choices in unequal current positions, 876 (97.8%) were in compliance with the intuitive rule. Interestingly, of the 20 noncompliant cases, all but two were from (3,4) or (0,1) pairings. Moreover, 7 of the 20 instances of non-compliance were the actions of a single subject.

In the case of players with equal current positions (the diagonal cells in Table 16), the equilibrium prediction of the model does less well, no matter what discount factor is assumed. There are too few observations to make much of the comparison with the theoretical calculations of p_{11} , p_{22} , and p_{33} from Table 2, but the extreme diagonal entries of 50.1% and 64.2% are lower and higher, respectively, than all the numbers in the p_{00} and p_{44} columns, respectively, for BOS(0) in Table 2. Indeed, the equilibrium predictions in these cells are evidently worse than the predictions of the one-shot mixed Nash equilibrium of the game, which entails \rightarrow play with probability .5. However, a t-test of the null hypothesis that the probability of right play is .5 when players of equal current positions are matched shows that even this hypothesis is rejected at conventional significance levels ($t = 2.1931$). This rejection is largely accounted for by the fact that (4, 4) matches result in too much \rightarrow play.

Table 17 reports the empirical frequencies of the positions for the last 50 rounds of BOS(0).

	Position				
	0	1	2	3	4
Frequency	.350	.093	.077	.097	.383

Table 17: Position Frequencies in BOS(0)

Notice that these empirical frequencies match up fairly closely with the predictions of the theoretical model in Table 5 for all discount factors.

BOS(5)

Table 18 presents the \rightarrow play data for the six BOS(5) sessions analogously to Table 16.

		Player's Own Current Position				
		0	1	2	3	4
Opponent's Current Position	0	50.4 (250)	87.9 (91)	91.4 (70)	94.5 (110)	97.3 (510)
	1	19.8 (91)	66.7 (18)	87.0 (23)	85.7 (35)	98.5 (130)
	2	8.6 (70)	39.1 (23)	78.6 (14)	96.0 (25)	96.7 (120)
	3	6.4 (110)	37.1 (35)	32.0 (25)	76.2 (42)	93.4 (168)
	4	2.2 (510)	10.8 (130)	14.2 (120)	41.1 (168)	67.8 (712)

Table 18: Play of \rightarrow in BOS(5)

Overall in the BOS(5) sessions, the compliance rate was 90.8% from unequal positions. Of the 235 instances of non-compliance, 80 were from (3,4) matches. The majority of these instances of noncompliance were, in turn, cases where in the previous period the subject had been in the extreme state 4 and played the nonaggressive strategy.¹⁹ Further, 93 of the 235 instances of non-compliance occurred in one session (#10). Overall, the compliance rates for BOS(5) seem high enough that any simple hypothesis about subjects play must involve something very close to the intuitive rule for the situations when the matched subjects are in unequal positions. As before, the vast majority of the 72 subjects in these sessions played either perfectly or almost perfectly in accord with the prescriptions of the intuitive rule. Specifically, only 10 subjects out of 72 had compliance rates below 85%.

From the equal-position pairs (0,0) and (1,1), \rightarrow was chosen far too infrequently relative to the equilibrium calculations of Table 3. Equilibrium predictions for (2,2), (3,3), and (4,4) for discount factors around .75 are more in line with the observed data. The one-shot mixed strategy equilibrium for this game is for both players to play \rightarrow with probability 62.5%, and the fact that the proportion of \rightarrow play from these three position pairs exceeded this level appears to reflect reputation-based reasoning.

Table 19 reports the empirical frequencies of each position observed in the last 50 rounds of the BOS(5) sessions. It seems closest to the BOS(5) π -distribution for the discount factor .85 in Table 2.

¹⁹This pattern recurs in other treatments. See Section 5 for attempts to explain it.

	Position				
	0	1	2	3	4
Frequency	.286	.083	.070	.106	.456

Table 19: Position Frequencies in BOS(5)

There is a pronounced rightward bias in both play and position frequencies in the BOS(5) data. This is due to two effects: (1) a bias toward the aggressive strategy when a player and the matched opponent are in the same position (which is qualitatively consistent with the predictions of the theoretical model in Section 2), and (2) a bias toward the aggressive strategy among the instances in which the prescription of the intuitive rule was not followed (about which the theoretical model is silent).

BOS(10)

Table 20 presents the \rightarrow play data for the two BOS(10) sessions. >From unequal positions, the compliance rate was 81.4%. As in the BOS(5) sessions, a significant proportion (60 out of 145) of the instances of noncompliance occurred in (3,4) matches. Excluding these pairings raises the compliance rate in BOS(10) sessions to 85.8%.

		Player's Own Current Position				
		0	1	2	3	4
Opponent's Current Position	0	59.4 (32)	92.9 (14)	93.8 (16)	90.9 (33)	93.6 (126)
	1	28.6 (14)	— (0)	75.0 (8)	100 (8)	91.9 (37)
	2	12.5 (16)	50.0 (8)	50.0 (2)	81.3 (16)	97.6 (42)
	3	6.1 (33)	50.0 (8)	50.0 (16)	86.7 (30)	85.4 (89)
	4	5.6 (126)	29.7 (37)	50.0 (42)	52.8 (89)	73.4 (358)

Table 20: Play of \rightarrow in BOS(10)

For players with equal current positions, most of the observations are (4,4) pairings, and the predictions from Table 4 are not too far off here (though they are for p_{00} , where far too little \rightarrow play is observed relative to the predictions of the theoretical model). In addition, because of the (4,4) data, we cannot reject the null hypothesis that equal-position players are simply choosing \rightarrow with the same probability as the one-shot mixed strategy equilibrium (probability=.75, $t = -0.8235$). It is hard to reconcile either of

these with the large number of noncompliant choices when 2's and 3's were matched with 4's.

The position frequencies are listed in Table 21.

	Position				
	0	1	2	3	4
Frequency	.184	.056	.070	.147	.543

Table 21: Position Frequencies in BOS(10)

Table 21 reveals that as in BOS(5), there is a pronounced rightward bias in play even relative to the π - distributions for BOS(10). This is evidently mostly due to the aforementioned noncompliant choices of 2's and 3's matched with 4's.

4.3 HD Results

We now turn to a detailed analysis of the results from the HD treatments.

HD(5)

Table 22 presents the \rightarrow play data for the two HD(5) sessions. From unequal positions, the compliance rate was 91.0%.

	Player's Own Current Position				
	0	1	2	3	4
Opponent's	0 27.1 (236)	58.5 (53)	96.9 (32)	93.3 (60)	98.1 (156)
Current	1 7.5 (53)	16.7 (6)	75.0 (8)	87.5 (8)	94.9 (39)
Position	2 9.4 (32)	12.5 (8)	50.0 (8)	71.4 (7)	93.1 (29)
	3 1.7 (60)	0.0 (8)	0.0 (7)	31.3 (16)	81.3 (32)
	4 3.8 (156)	12.8 (39)	24.1 (29)	12.5 (32)	27.9 (26)

Table 22: Play of \rightarrow in HD(5)

For this treatment, the bulk of non-compliant play did not come from (3,4) pairings: only 10 of the 76 instances of non-compliance were from (3,4) matches. Instead, the largest portion (26) of the 76 non-compliant plays came from (1,0) matches. This change from the BOS pattern is due, no doubt, at least in part, to the fact that in HD failing to coordinate by playing \leftarrow is less costly than failing to coordinate by playing \rightarrow . Although the

equilibrium calculations take account of this effect, our subjects evidently overcompensated: as evidenced by much-too-low (relative to equilibrium) incidence of \rightarrow play from equal positions. Notice that in HD(5) sessions, the mixed-strategy equilibrium prediction of the one-shot game is that \rightarrow will arise with probability .25. In a t-test, we fail to reject the null hypothesis that the proportion of \rightarrow choices from equal positions is equal to the one-shot mixed strategy probability ($t = 1.1875$) at conventional significance levels.

The position frequencies are listed in Table 23.

	Position				
	0	1	2	3	4
Frequency	.448	.095	.070	.103	.285

Table 23: Position Frequencies in HD(5)

These empirical frequencies are closest to the theoretical π - distribution for relatively low discount factors, but, unlike the BOS treatments, the empirical position frequencies in HD(5) reflect a leftward bias even relative to the left-biased predictions of the theoretical model for low discount factors. Even for discount rates as low as $\beta = .5$, there are more 0 types and fewer 4 types than is predicted by the equilibrium corresponding to the intuitive rule. This may all be accounted for by the bias toward \leftarrow play in equal-position matches, but the underlying cause may be other-regarding preferences, such as a taste for fairness, that might encourage subjects to try to coordinate on (\leftarrow, \leftarrow) play; thus achieving an equal split of the available earnings in a given round.

HD(10)

Finally, we come to the four HD(10) sessions; Table 24 presents the \rightarrow play data. From unequal positions the overall compliance rate was only 72.9%.

	Player's Own Current Position				
	0	1	2	3	4
Opponent's	0 27.4 (190)	56.7 (97)	62.5 (96)	75.0 (100)	90.8 (238)
Current	1 15.5 (97)	37.0 (46)	40.0 (30)	62.2 (37)	86.1 (101)
Position	2 20.8 (96)	33.3 (30)	35.3 (34)	56.0 (25)	85.7 (84)
	3 31.0 (100)	37.8 (37)	40.0 (25)	71.4 (28)	74.5 (106)
	4 20.6 (238)	48.5 (101)	36.9 (84)	42.5 (106)	64.2 (274)

Table 24: Play of \rightarrow in HD(10)

As the table makes clear, the predictions of the intuitive-rule equilibrium did not fare at all well when players in positions 1-3 were matched with one another. Players with current position 0 and 4 played in a manner that is more consistent with the intuitive rule. If we pool all of the observations of choices by players from equal positions and compare this to the proportion of \rightarrow choices predicted in the one-shot mixed-strategy equilibrium (.5), we again cannot reject the null hypothesis at conventional significance levels. However, this is largely due to the fact that the prevalence of \leftarrow play from positions (0,0), (1,1) and (2,2) is being offset by the prevalence of \rightarrow play by (3,3) and (4,4) pairs rather than the fact that players with equal current positions are playing the one-shot mixed strategy.

The position frequencies are listed in Table 25.

	Position				
	0	1	2	3	4
Frequency	.300	.130	.112	.123	.335

Table 25: Position Frequencies in HD(10)

Relative to the equilibrium calculations, position 4 is underrepresented. This leftward bias appears less pronounced than in the HD(5) sessions however.

5 Discussion

The experimental results appear to us to support the following:

1. Subjects attempted to use position pairs as coordinating devices eventually in all sessions, and the intuitive rule, or something close to it, was the custom that invariably arose.
2. Adherence to the rule decreased in the $y-$ parameter for each game form.
3. Aggregate play in each session settled down to the semblance of a stationary pattern by Round 50.
4. Adherence to the play predictions of the steady state equilibrium corresponding to the intuitive rule was better from unequal position pairs than from equal position pairs.

We repeat our earlier caveat that little of this can be established statistically both because the intuitive rule admits no variance from unequal

position pairs and because we have not run nearly enough sessions of any one treatment.²⁰ Furthermore, it could be that behavior depends in unexpected ways on the specific treatment variables that we kept fixed. There is an unlimited number of variations that would be interesting to try: total number of rounds fixed $\neq 100$, random number of rounds not told to subjects in advance, diminishing payoff magnitudes over time to simulate discounting, different x and y parameters, number of positions $\neq 5$, current position distribution hidden from subjects, different monetary scale, different subject pool, etc. It is impossible to know whether any of these variations would produce dramatically different patterns.

Assuming that we are seeing a pattern that would be robust to such variations, does it occur because subjects are conditioned to the intuitive rule from their ordinary life experiences? Assuming yes to this, is the intuitive rule so pervasive in the real world because of the coordination role analogous to what is isolated in this experiment, or could it be that it serves a variety of functions of which this is only one? It is impossible to know, but we suspect the coordinating role is at least one of its important functions if not its most important function. This leaves the question of why the intuitive rule is so prevalent when there are so many other possible coordinating rules. Despite the suggestions in Section 1, we do not really know.

Although there are obvious mismatches between the assumptions of the model in Section 3 and the experiment, the model does reasonably well at matching up with the empirical π -distribution, although less well for the $\{p_{ii}\}$. This should perhaps not be too surprising given the number of experiments on one-shot games in which subjects do not appear to play mixed strategies in equilibrium proportions.²¹

Of course, with subjects being paid only at the end of the experiment, it is not obvious why they should discount at all, let alone geometrically or with particular common discount factors. While we cannot defend as realistic the specific form of the discounting assumed in the model, we suspect that subjects do implicitly trade off present against future winnings as they make choices in each round. And, perhaps because they are accustomed to accept unthinkingly the advice about the relative value of a bird in the hand, or because they are risk averse and perceive less risk in the results of play in

²⁰One could argue that because of the degree of interaction among subjects in a given session, one session might properly be treated as a single observation.

²¹See, e.g., Brown and Rosenthal (1990), Rapoport and Boebel (1992), Mookherjee and Sopher (1994), Shachat (1996).

the current round, or for some other reason, for the most part they seem to discount the future. Evidence for this is, somewhat circularly, based on the observation that they mostly play according to the intuitive rule. Consider for instance, a single individual in a large population of people playing BOS with positive y according to the intuitive equilibrium for a particular common discount factor. Since v_4 is the highest equilibrium present value, if the individual's discount factor were higher than that of the rest of the population or if the individual didn't discount at all, the individual would not be indifferent in the equal-position matches, would eventually get to position 4, and would always play to the right. A small proportion of our subjects did seem to behave this way, but only a small proportion. If most subjects did, the average frequency of position 4 would be high and so would the instances of $(\rightarrow, \rightarrow)$ play (which gives a hint about why existence fails in the model when the discount factor is too high relative to y). This also suggests why the theoretical model may not predict too badly if most subjects discount the future, even if they do not all have a common discount factor or do not discount geometrically. Because of the anonymity and the relatively large numbers of subjects in each session, what is important is that there is enough discounting going on in the aggregate. Even if nobody is exactly indifferent at any equal-position match, if there is enough heterogeneity each individual faces an environment that looks as though all opponents are behaving similarly, though not deterministically. And it is only the conjectures individuals have about their aggregate environment that matter in their respective optimization problems. In other words, there are analogues of steady state equilibria for repeated anonymous matching games played by heterogeneous populations, but, when they exist, their aggregate behavior is always akin to that of equilibria of some homogeneous population model.

Similar observations can be used to rationalize the use of an infinite-population model when the population size is in fact finite. For some large enough number of subjects (we hope 12 is large enough), with anonymity reasonable individuals should be playing as if there is some undifferentiated mass of opponents.

As for the assumption that the stage-game money payoffs are von-Neumann-Morgenstern utilities, since the money payoff in any round is small compared with total money payoff, risk neutrality ought not to be too far wrong.

The apparent absence of end-game effects in the data, justifying the convenient use of a stationary model, is not surprising given the nature of the stage games in this experiment. In the unequal position matches, which form

the large majority of matches in the data in most rounds, the intuitive rule prescribes play of a one-shot Nash equilibrium. Thus, if an individual in such a match is convinced that his opponent will play according to the rule in the last round, there is no reason for the individual to deviate in the last round. In the penultimate round, if the individual does not expect a last-round position distribution that is much different than that of the steady state equilibrium and if his matched opponent occupies a different position than his own, similar considerations apply. For equal-position matches in the last round, theoretically there is the symmetric mixed equilibrium of the stage game, but we would hesitate to rely on that (or anything else) for a prediction. In fact, for the purpose of explaining the absence of end-game effects what is most important is only that whatever the subjects themselves predict about equal-position matched play in the last round does not influence incentives in the penultimate round when the subjects are matched with an opponent in an unequal position.

Our explanation for the monotonicity of compliance rates in y in the data is related to the discussion of discounting above. We suspect that if the experiment could be controlled so that the parameters and discount factors were such that the intuitive rule did not form an equilibrium, our subjects would be at sea. Perhaps after a much longer play sequence they would manage to settle on one of the rules that do form equilibria, but this would not be something that their real-world experiences would help with. And in rounds 51-100 of such a treatment we have no idea what would happen. So, consider the thought experiment in which it is common knowledge that all players discount geometrically, but that each player's discount factor is a random draw from some distribution of discount factors, with positive probability on the event that the aggregate amount of discounting is insufficient for the intuitive rule to form an equilibrium (of the heterogeneous-population model with complete information about all discount factors). Our guess is that in such an experiment we would see some subjects attempting to play the rule and some subjects groping about; the lower the y -value the less groping and the higher the compliance rate. In this sense then, the model (or at least some imagined variant of it) could explain the observed monotonicity in y .

Might there be alternative explanations for monotonicity of compliance rates in y ? Perhaps. First, there is intuition for the idea that the higher is y the more one is willing to risk the worst outcome for a chance at the best outcome. So, if an individual thinks that the matched opponent might play the intuitive rule with some error, the individual might be more likely

to try to depart when y is high than when it is low (although the necessary assumptions on the error-generating process might not be realistic). But it seems possible to us that this could be formalized anyway with a generalized version of the theoretical model of Section 3 and so is not necessarily an alternative to it. Second, if y is high and an individual thinks the intuitive rule is being played, the individual may choose to sacrifice current payoff when in a lower position than that of the matched opponent by departing from the rule in order to get to a higher present-value position. But this is characteristic behavior of a high-discount factor player, independently of y . Such players will soon find themselves in the extreme state and will be in compliance with the intuitive rule after Round 50.

We noted earlier the pattern of departures from the intuitive rule in BOS(5) and BOS(10) in which a 4 meets a 4, plays \leftarrow , then plays \rightarrow against a 4 in the next round. In the simplistic terms of our discussion about discounting, it is as though the player first acted as though she had a discount factor no higher than the common one, then, in the next period had one that was higher. While we cannot rationalize this within the confines of our model, it does strike us as reminiscent of behavior in other experiments (see Ainslie and Haslam, 1992, for a survey) that is more consistent with a hypothesis of hyperbolic discounting, and hence time-inconsistency, than with geometric discounting.

Finally, we point out one of the other steady state equilibria of the model that has features that are reflected in some of the data. It is generated by a variant of the intuitive rule in which individuals act as though positions 3 and 4 are the same. In other words, when an individual in position 3 meets an individual in position 4, the first randomizes as though he were in position 4 and the second randomizes as though she were in position 3. It therefore has a chance to provide an explanation for the pattern of departures from compliance discussed in the previous paragraph. As an example, for BOS(5), such steady state equilibria exist for $\beta < .95$. For instance, when $\beta = .9$,

$$p_{00} = .78, p_{11} = .93, p_{22} = .83, p_{33} = p_{43} = .69, p_{44} = p_{34} = .64.$$

$$\pi_0 = .27, \pi_1 = .09, \pi_2 = .06, \pi_3 = .12, \pi_4 = .46.$$

This does well at explaining the position frequencies, but less well at explaining the lack of \rightarrow play when pairs with equal current positions are matched.

There are many other hybrids.

A Appendix: Instructions to Subjects

This is an experiment in decision-making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money. You will be paid in private and in cash at the end of the session.

There are twelve people participating in this session. It is important that you remain silent once play has begun. If you do not, the experiment will be terminated, and you will receive none of your prize money. If you have a question during the experiment, raise your hand and a monitor will come over to where you are sitting and answer your question in private.

The experiment will consist of a simple game played 100 times in succession. In each of the 100 rounds you will be asked to click either the left arrow (\leftarrow) or right arrow (\rightarrow) on the screen with your mouse. Your winnings for the round will be 5¢, 20¢, or 30¢, depending on the choice you make and the choice made by your co-player for that round. Thus your total winnings for the experiment will be between \$5.00 and \$30.00. Since you will be paid in cash, privately, at the end of the experiment, none of the other participants in the experiment will know what your winnings are. You will also never be identified by name on any publications or reports on the results from this experiment.


In each round, your co-player will be one of the other experimental subjects in this room. They have all been recruited as subjects in the same way that you have, and they are reading exactly the same instructions that you are for the first time now. In each round your co-player will be chosen anew according to a completely random drawing. Thus, in each round your co-player is equally likely to be any of the other subjects, independently of anything that has transpired before that round. Neither you nor your co-player will be told the identity of the other.

Winnings for you and your co-player in each round are determined as follows: If both of you choose the right arrow (\rightarrow), you both get 5¢. If both of you choose the left arrow (\leftarrow), you both get 5¢. If one of you chooses (\leftarrow) and the other chooses (\rightarrow), the one who chose (\rightarrow) gets 30¢ and the other gets 20¢. To summarize, your payoff table is

		Co-Player's Choice	
		\leftarrow	\rightarrow
Own Choice	\leftarrow	5¢	20¢
	\rightarrow	30¢	5¢


All players have exactly the same payoff table.

Although the identities of matched co-players will not be revealed in any round, each of you will receive a piece of information about the other's recent play history before choosing which arrow to click. This piece of information is called the player's current position, which can be at any of the five locations on the line segment below.




A player's current position is normally arrived at by taking his/her previous position and shifting it one location to the right if the player's previous choice was (\rightarrow) or one location to the left if the previous choice was (\leftarrow).

However, if the previous position was



and the player chose (\rightarrow), the new position will simply be the same as the previous position. Similarly, if the previous position was




and (\leftarrow) was chosen, the position does not change.

You will be reminded in each round as to your own current position as that is information that your co-player has about you. Each player will also be kept informed of the distribution of ALL players' positions. In round one, for instance, all 12 players will begin at the middle position, so the initial screen display will be:


Round 1 of 100

Total Earnings so far \$0


OWN Current position



CO-PLAYER'S current position



Current position frequencies



Payoff Table:

		Co-Player's Choice	
		←	→
Own Choice	←	5¢	20¢
	→	30¢	5¢

Please Choose an arrow

← →

When all players in the room have made their choices in a round, your payoff for the round will be displayed in the form of a highlighted cell in the payoff table. To begin the next round, simply click the mouse anywhere on the screen.


Are there any questions?

References

- [1] Ainslie, G. and N. Haslam, “Hyperbolic Discounting”, in *Choice over Time*, G. Loewenstein, ed., Russell Sage Foundation, New York, 1992.
- [2] Brown, J. and R. Rosenthal, “Testing the Minimax Hypothesis: A Re-examination of O’Neill’s Game Experiment”, *Econometrica*, 58 (1990), 1007-1028.
- [3] Cooper, R., R. Forsythe, D. DeJong, and T. Ross, “Alternative Institutions for Resolving Coordination Problems: Experimental Evidence on Forward Induction and Preplay Communication,” in *Coordination in Economic Activity*, J. Friedman, ed., Kluwer Academic Press: Amsterdam, 1994.
- [4] Mookherjee, D. and B. Sopher, “Learning Behavior in an Experimental Matching Pennies Game”, *Games and Economic Behavior*; 7 (1994), 62-91.
- [5] Rapoport, A. and R. Boebel, “Mixed Strategies in Strictly Competitive Games: A Further Test of the Minimax Hypothesis”, *Games and Economic Behavior*; 4 (1992), 261-83.
- [6] Rosenthal, R. and H. Landau, “A Game-Theoretic Analysis of Bargaining with Reputations”, *Journal of Mathematical Psychology*, 20 (1979), 233-255.
- [7] Roth, A. and M. Malouf, “Game-Theoretic Models and the Role of Information in Bargaining”, *Psychological Review*, 86 (1979), 574-594.
- [8] Schelling, T., *The Strategy of Conflict*, Harvard university Press, Cambridge, 1960.
- [9] Selten, R., A. Sadrieh, and K. Abbink, “Money Does Not Induce Risk Neutral Behavior but Binary lotteries Do Even Worse”, *Universitat Bonn Sonderforschungsbereich 303, Discussion Paper B 343*, 1995.
- [10] Shachat, J., “Mixed Strategy Play and the Minimax Hypothesis”, *UCSD Economics Discussion Paper 96-37*, 1996.

Round 1 of 100

Total Earnings so far \$0



Payoff Table

		co-choice	
		←	→
own choice	←	25¢	15¢
	→	35¢	5¢

Please choose an arrow.




Figure 1: Sample Screen Image




Figure 2: Average compliance rates by round -- all treatments




Figure 3: Average compliance rates by round – BOS(0)




Figure 4: Average compliance rates by round – BOS(5)




Figure 5: Average compliance rates by round – BOS(10)




Figure 6: Average compliance rates by round – HD(5)




Figure 7: Average compliance rates by round – HD(10)