

Adaptive Learning versus Punishment in Ultimatum Bargaining¹

Klaus Abbink

Universitat Pompeu Fabra, and Institut d'Anàlisi Econòmica CSIC, Barcelona, Spain

Gary E. Bolton

Smeal College of Business, Penn State University, University Park, Pennsylvania 16870

Abdolkarim Sadrieh

*Department of Economics, Tilburg University, PO Box 90153, 5000 LE Tilburg,
The Netherlands*

and

Fang-Fang Tang

Nanyang Business School, Nanyang Technological University, Singapore 639798

Received August 20, 1998; published online August 8, 2001

Adaptive learning and a fairness motive we call “punishment” are the basis for two prominent and substantially different types of theories of ultimatum bargaining behavior. We compare adaptive learning and fairness in an experiment that involves punishment and reward versions of the ultimatum game. We draw conclusions concerning the abilities and limitations of both types of theories. The results shed light on how learning and fairness interact, information that should be useful in constructing a more comprehensive model. *Journal of Economic Literature*
Classification Numbers: C78, C91, D83. © 2001 Academic Press

¹Special thanks are extended to Reinhard Selten for his advice and encouragement. Funding for the project was provided by the Deutsche Forschungsgemeinschaft through the Sonderforschungsbereich 303. This project began while Bolton was a visitor at the Institut für Gesellschafts- und Wirtschaftswissenschaften, Bonn. He gratefully acknowledges the Institut’s hospitality and financial support. Bolton also gratefully acknowledges the support of the National Science Foundation.

1. INTRODUCTION: TWO INTERPRETATIONS OF ULTIMATUM GAME PLAY

The ultimatum game has become synonymous with debate over the nature of strategic reasoning and the role fairness plays in bargaining. The debate has progressed markedly as evidence has accumulated and new theories to account for that evidence have come into being. Two different types of theories, one centered on a fairness motive and the other centered on adaptive learning, dominate the recent literature. Along some lines the theories clearly compete; along others they are conceivably complementary. The precise nature of the relationship is cloudy in part because each theory, by itself, is in basic accord with existing data. We report here on a new experiment designed to clarify the relationship.

In the ultimatum game, a first mover proposes a division of a fixed monetary sum to a second mover. If the second mover accepts, the money is divided accordingly; if he or she rejects, both players receive nothing. Conventional subgame perfect equilibrium predicts that the first mover will receive virtually the entire sum. Nevertheless, beginning with Güth *et al.* (1982), laboratory investigators have found that first movers tend to offer amounts that are substantially higher than the minimum, and second movers tend to reject if accepting means taking a relatively small share (Roth, 1995).

The *punishment hypothesis* asserts that, along with the monetary payoff, some bargainers care about how fair the division is to self. While there are several versions of this argument, the comparative bargaining model (Bolton, 1991) provides a simple formulation that captures the essentials. The model posits that bargainers care about relative, as well as pecuniary, payoffs. Second movers then reject inequitable offers in order to obtain a more favorable relative division. (The term “punishment” serves to distinguish this explanation from one positing that first movers want to make fair offers.) Bolton shows that the model can explain various observations, mostly comparative statics, from two-period alternating offer bargaining experiments. Application to the simpler ultimatum game is immediate. The model has since been extended to other games.²

The *adaptive learning hypothesis* posits that ultimatum bargainers modify their behavior on the basis of the outcomes they experience in previous

² Rabin’s (1993) concept of fairness equilibrium leads to a version of the punishment hypothesis based on the idea that people want to hurt people who hurt them, a notion that is different than relative payoffs. Blount (1995) and Kagel *et al.* (1996) offer empirical comparisons of the two conceptualizations. Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) demonstrate that generalized versions of the relative payoff model can explain lab observations from a variety of bargaining, dilemma, and market games. Camerer (1997) reviews the development of earlier fairness models.

play. The model of this kind most closely associated with the ultimatum game is reinforcement learning (Roth and Erev, 1995). The model takes an initial propensity to reject some offers as given and then simulates the path of repeated play. The better the monetary payoff from an action taken, the more likely the action is repeated in the future. Roth and Erev show that when initial playing propensities are fit from early round data, reinforcement learning simulations track the path of ultimatum play. Adaptive learning has been applied to explain a variety of other games.³

Bolton (1991) informally discusses learning, but only with respect to changes in first mover behavior. Roth and Erev (1995) do not discuss fairness, but say “[t]hat such simple dynamic models, when initiated with first round observed behavior, nevertheless do a good job of predicting how observed behavior will evolve, suggests that a substantial part of how players’ knowledge and beliefs influence the game may be reflected already in first round data” (p. 204).

There is overlap in what fairness and adaptive learning set out to explain. Of particular importance, both models characterize second mover rejections (in quite different ways). But the overlap is not exact. Fairness models do not address learning, and adaptive learning models do not model strategic thinking. The fundamental difficulty in testing one model against the other, or in determining the relationship between them, is that fairness models have been formally worked out in static equilibrium frameworks with the aim of predicting comparative statics, whereas adaptive learning is inherently dynamic and focused on explaining the influence of experience.

Our experiment makes use of a version of the ultimatum game in which the first mover does not know whether he or she receives a “reward” payoff or “punishment” payoff when the unequal split is rejected, but the second mover does know. It turns out that, for this game, what second movers know about first mover payoffs is irrelevant to what the adaptive learning hypothesis predicts—at least so long as we suppose no difference in second mover initial playing propensities across information realizations (a nonlearning explanation). In contrast, this information is of crucial importance to what the punishment hypothesis predicts.

³ Reinforcement learning has many features in common with a class of adaptive learning models introduced by Bush and Mosteller (1955); see Tang (1996). Harley (1981) proposed a model of this type as an approximation for evolutionary dynamics. Roth and Erev (1995) also examine the best shot game and a simple auction market. Chen and Tang (1998) examine a public goods game, Duffy and Feltovich (1999) do a further analysis of ultimatum and best shot games, Erev and Rapoport (1998) analyze a market entry game, and Erev and Roth (1998) examine games with unique mixed strategy equilibria. Gale *et al.* (1995) apply “noisy” replicator dynamics to the ultimatum game as an evolutionary explanation for the survival of fairness norm.

The main treatment effect of our experiment, that second mover rejection rates are higher in the punishment treatment than the reward treatment, suggests that the punishment hypothesis offers a more accurate explanation of rejection behavior than does (slow) adaptive learning. On the other hand, we find evidence for adaptive learning among first movers. A natural reaction to these results is to wonder whether an adaptive learning model that fits initial second mover propensities conditional on first mover payoff information might provide a comprehensive explanation. We examine a simple adaptive learning model similar to the reinforcement learning model. The simulations track the treatment with the usual ultimatum payoffs but fail to do so in the award payoffs treatment where rejection rates are lower than what the model would predict.

The results suggest that the role fairness plays in the ultimatum game goes beyond initial propensities. We think this is important information for future theory building. A more comprehensive model—one that robustly describes comparative statics as well as tracks dynamics—will have to explicitly grapple with how fairness and learning interrelate during repeated play.

2. DESIGNING A SEPARATING EXPERIMENT: A CLOSER LOOK AT HOW THE HYPOTHESES DIFFER

While conventional perfect equilibrium does not accurately forecast ultimatum play, it does provide a useful framework for thinking about the differences between adaptive learning and punishment explanations. For the ultimatum game, perfect equilibrium makes three assumptions:

P1: Each bargainer's exclusive motivation is to gain the most money possible.

P2: Each bargainer knows the other bargainers' motivation.

P3: Each bargainer can identify his or her optimal action.

The equilibrium prediction is then constructed as follows: Since the first mover knows that the second mover prefers more money to less (P1 and P2), the first mover should offer the second mover the smallest monetary unit allowed, allocating the balance to himself (P3). The second mover should accept (P1 and P3). We can think of the punishment and adaptive learning hypotheses as amending different postulates of the conventional perfect equilibrium argument.

2.1. *The Punishment Hypothesis*

The punishment hypothesis attributes a motive to second mover rejections, specifically, to mitigate the unfairness of an unequal division. First movers then tend to shy away from the perfect equilibrium offer out of fear of winding up with nothing.

Bolton (1991) expresses the motive in terms of bargainer preferences. In the context of the ultimatum game, bargainer n 's utility function over settlements is taken to be

$$u^n(x, i(x, z)), \quad (1)$$

where x is the monetary payoff to n and z is the monetary payoff to his or her partner, $i(x, z) = x/z$ is the relative measure of the settlement, with $i(0, 0) = 1$, $u_1^n > 0$, $u_2^n > 0$ when $i < 1$, and $u_2^n = 0$ otherwise. An ultimatum game second mover is then better off turning down any offer that gives him or her utility less than $u^n(0, 1)$, the amount he or she obtains by rejecting.

In term of the perfect equilibrium postulates, the punishment hypothesis maintains P2 and P3 but replaces P1 with

P1*: For divisions that offer him or her a sufficiently small share of the settlement, the bargainer prefers that both receive nothing; otherwise, he or she prefers the division that offers the most money.

The Eq. (1) preferences are one example of P1*. The Bolton (1991) model exhibits perfect equilibria in which relatively small offers will be rejected, and agreed upon settlements are somewhere between equal division and the conventional perfect equilibrium. Bolton's (1991) model assumes complete information, but Bolton and Ockenfels (2000) extend the analysis to incomplete information.

2.2. *The Adaptive Learning Hypothesis*

Perfect equilibrium analysis is static. Adaptive learning focuses on dynamics. Succinctly put, the adaptive learning hypothesis posits that second movers learn to accept more generous offers more quickly than less generous ones. This pushes first movers to more generous offers.

Roth and Erev's (1995) reinforcement model takes as given that each bargainer n begins the first round, $t = 1$, with an *initial propensity* to play his or her k th pure strategy, given by some number $q_{nk}(1)$. The probability that bargainer n will play k in round 1 is $q_{nk}(1)/\sum_j q_{nj}(1)$, where the denominator is known as the *initial strength*. Repeated play modifies propensities through a process of adaptation. If n plays his or her k th pure

strategy in round t and receives the payoff x , then the propensity to play k is updated to

$$q_{nk}(t + 1) = q_{nk}(t) + x. \quad (2)$$

The probability that k is played in round $t + 1$ is then tabulated as it is in round 1.

The model holds to strict self-interest (P1), if only implicitly: bargainers learn in the direction of more money for self. But the model replaces P2 and P3 with

P2*: Each bargainer has some initial propensity to play his or her k th pure strategy, $q_{nk}(1)$.

P3*: Each bargainer adapts propensities as specified by Eq. (2).

P2* allows second movers to have an initial propensity to reject offers and leave money on the table. The considerations behind this propensity are not modeled. One possibility, consistent with the other assumptions made, is cognitive mistakes.

Reinforcement model predictions are derived from computer simulations, first choosing an initial strength level and fitting initial propensities from early round data. Roth and Erev find that predicted ultimatum game behavior settles down away from conventional perfect equilibrium for long periods. Gale *et al.* (1995) obtain similar results for the miniultimatum game using an adaptive learning algorithm derived from replicator dynamics.

3. GAME WITH AN UNCERTAIN REWARD: THE NEW EXPERIMENT

The new experiment involves a version of the cardinal (or mini) ultimatum game in which the first mover has exactly two offers to choose from. The nature of behavior we observe in the cardinal ultimatum is much the same as in the standard ultimatum game (see Bolton and Zwick, 1995). The advantage of cardinal ultimatum is that it simplifies the experimenter's measurement task by providing an unambiguous benchmark of "success." Specifically, our experiment separates the two hypotheses by detecting differences in the rate of perfect equilibrium responses. If we were to allow many offers there would be ambiguity as to whether some observations were "close" to perfect equilibrium. For cardinal ultimatum, the judgement is clear-cut.

3.1. A Separating Test That Introduces Incomplete Information about the Payoffs

The game is presented to players as displayed in Fig. 1, with “?” standing in for the first mover’s payoff when the second mover rejects the unequal split. Prior to play, it is publicly announced that ? has been randomly selected to be either 0 or 10, each having equal chance; the value is the same for all games in the session (given that each treatment would have a total of five sessions). The second mover is privately told the value of ?; that the second mover has this information is known to all players. First movers play a series of single shot games with second movers; no two bargainers matched more than once. Players observe the actions of their partners but not the actions of non-partners. The value of ? is not revealed to first movers until all games are complete. Even if the first mover receives the ? payoff, he or she must wait until the end of the session to find out its value.

It is important to understand the special nature of the incomplete information in this game. The first mover has incomplete information about his or her *own* payoff, but the second mover knows all payoffs for certain. This particular incomplete information is of no substantive significance to the perfect equilibrium analysis: the value of ? does not affect the second mover’s payoffs and so should not affect the second mover’s actions. The first mover, therefore, can expect to get the unequal split if he or she moves right, independent of the value of ?.

More important to our purposes, the realization of the value of ? is irrelevant to the adaptive learning hypothesis—at least so long as we do not invoke a nonlearning explanation such as differences in initial propen-

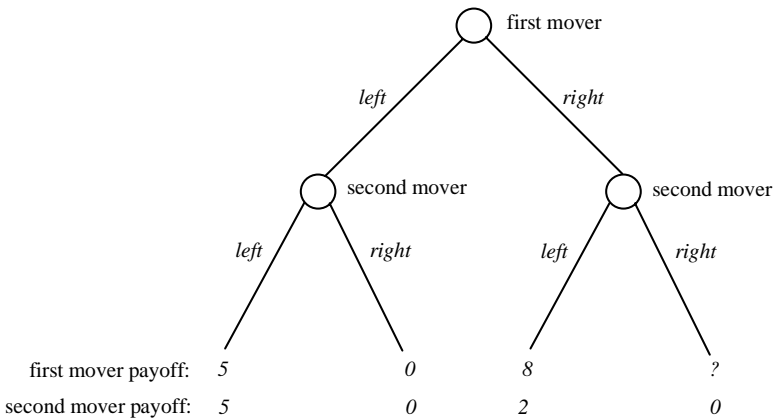


FIG. 1. The uncertain reward version of cardinal ultimatum.

sities across treatments (a restriction we will relax in Section 5). Start with the first round of the round robin. Since the game looks identical to first movers for both values of β , there should be no systematic difference in first round, first mover offers. Supposing no difference in initial propensities due to differences in first mover payoffs (a nonlearning argument), adaptive learning implies that the value of β should have no influence on second mover behavior. And since second movers will be responding to essentially the same set of first mover offers, there should be no systematic difference in second mover responses. Recall that adaptive learning posits that players modify their behavior in response to their own past payoffs. With no systematic difference between first round histories, there should be no systematic difference between second round, first mover offers and no systematic difference in second mover responses, etc. In sum, there should be no systematic difference across treatments. Accordingly, our experiment places adaptive learning in the position of the null hypothesis.

In contrast, the punishment hypothesis predicts a sharp difference. When $\beta = 10$, moving right in response to the unequal offer effectively rewards the first mover, rather than punishing him or her as it does when $\beta = 0$. Put another way, when $\beta = 10$, the second mover is better off playing accept in both pecuniary and relative terms. Hence the punishment hypothesis is the alternative hypothesis, predicting that, when $\beta = 10$, second movers will exhibit a lower propensity to reject unequal offers and that the proportion of games ending in perfect equilibrium will be higher.

3.2. *Procedures, Subject Pool, and Sample Sizes*

The experiment was run in the Laboratorium für experimentelle Wirtschaftsforschung at the University of Bonn using RatImage software (Abbink and Sadrieh, 1995). The experiment involved a total of 160 subjects, all recruited through billboards posted around campus. Most were law or economics students. Participation required appearing at a special place and time, and was restricted to one session. The chance to earn cash was the only incentive offered.

Each treatment was comprised of five sessions. The order of sessions, given in Appendix B, was selected by random draw. In each session, eight subjects were randomly assigned the role of first mover, and eight the role of second mover. Each first mover played each second mover in random order over eight rounds.

At the beginning of each session, subjects read instruction sheets (copy in Appendix A). A monitor orally reviewed the procedures. Each subject was then seated in an isolation cubicle and played the game via computer, each time with a different, anonymous partner. The game appeared on the screen in decision tree format, essentially as it is presented in Fig. 1.

Second mover screens displayed the value of β , but first mover screens did not. Payoffs appeared on the screen as “talers,” in the same quantities displayed in Fig. 1, with β equal to 0 talers in the punishment treatment and equal to 10 talers in the award treatment. At the end of the session, talers were redeemed in cash at a rate of 50 pfennings per taler.⁴

4. TESTING THE PREDICTIONS

In this section, we test the predictions developed in Section 3. We find that the data deviate from the adaptive learning hypothesis in ways that the punishment hypothesis predicts. But we also find evidence for first mover learning. In the next section, we study whether the deviations from the adaptive learning hypothesis can be accounted for by permitting different initial propensities for the two treatments, with adaptive learning explaining the rest.

The entire data set appears in Appendix B. We have five sessions—five fully independent observations—for each treatment.

Second mover response to an offer of an unequal split is the key difference in how adaptive learning and punishment apply to the experiment. Here are the percentage rates at which second movers rejected unequal splits, by session (ordered):

$$\begin{aligned} \beta = 10: & \quad 0.0 \quad 1.3 \quad 3.5 \quad 12.0 \quad 12.3; & \quad \text{mean} = 5.8 \\ \beta = 0: & \quad 3.4 \quad 16.1 \quad 25.0 \quad 25.5 \quad 27.8; & \quad \text{mean} = 19.6. \end{aligned}$$

The average rate of rejection for $\beta = 0$ is in line with that typically observed in a regular ultimatum game (15–20%; Roth, 1995). The rate for $\beta = 10$ is less than a third as large as for $\beta = 0$. A permutation test rejects the hypothesis that the five-session average is the same across treatments in favor of the hypothesis that it is lower in $\beta = 10$ (one-tailed $p = 0.020$).⁵

Figure 2 provides a second perspective on rejecting behavior. The cumulative distribution of second mover rejection rates for $\beta = 10$ dominates that for $\beta = 0$. Only 5% of $\beta = 10$ second movers reject more than 20% of the unequal-split offers they receive, whereas 24% of $\beta = 0$ second movers do so. The $\beta = 0$ second movers are more likely to have a positive

⁴ Subjects were paid solely what they earned for the games. Average payment per subject was DM18.44, approximately \$14.75 at the exchange rate of the time. On average, each session took 45 minutes.

⁵ The p -value for the permutation test is the probability of getting a difference in average as extreme as the one observed when one randomly divides the ten observations into two equal sized groups. The permutation test is also known as the Fisher randomization test; see Davis and Holt (1993, p. 542–544) or Moir (1998).




FIG. 2. Rate at which second movers reject the unequal split, cumulative distribution.

rejection rate (z -test one-tail $p = 0.046$) and are more likely to have a higher average rate of rejection (z -test one-tailed $p = 0.013$). Only 2 of 9 of the $? = 10$ rejecters have multiple rejections, whereas a majority of $? = 0$ rejecters, 9 of 16, have multiple rejections.

The lesser rate of rejection in the $? = 10$ treatment is in line with what we would expect from the punishment hypothesis. However, there are some rejections in four of the five $? = 10$ sessions; the punishment hypothesis anticipates none. As we will see below, rejections in $? = 10$ sessions tend to be more prevalent early on, tailing off to zero by the last round. A potential explanation for this is that second movers reject the unequal split in early rounds in order to induce first movers to offer the equal split in future rounds. It is, however, difficult to reconcile this story with conventional reputation-building models (e.g., Kreps *et al.*, 1982). These models rely on the same players having repeated interaction, whereas our subjects were never matched together for more than one round (and were told such prior to play; see Appendix A).

Further evidence for separation comes from an examination of round-by-round differences in games played in (subgame) perfect equilibrium. We see from Fig. 3 that, although $? = 10$ dominates $? = 0$ in every round, the difference in perfect equilibria is more pronounced in the later rounds. Comparing the proportion of perfect equilibria in round 1, we find no evidence for a difference across treatments in perfect equilibrium play (sample of 40 per treatment, proportions of 0.65 and 0.60, respectively, one-tailed $p = 0.322$). But repeating the same test on the last round of data, we can reject the null hypothesis in favor of the hypothesis that the proportion of perfect equilibria is greater in $? = 10$ (proportions of 0.83 and 0.63, respectively, one-tailed $p = 0.023$).




FIG. 3. Games played in perfect equilibrium (each treatment: five sessions \times eight games/round).

So we have two types of evidence consistent with the comparative static predictions of the punishment hypothesis. These are the difference in second mover rejection rates of the unequal split and the late round difference in the proportion of perfect equilibria.

In Fig. 4, we see that rejections in $? = 0$ exhibit no trend.⁶ While second mover rejections in $? = 10$ trail off in the later rounds, neither simple inferential tests nor probit analysis turn up any statistical evidence for second mover learning.

One explanation why differences in perfect equilibrium play show up only in later rounds is that first movers in the two treatments learn different things about second mover behavior. In fact, in Fig. 4 we see that first mover behavior separates across treatments after round 3 in much the way we would expect if first movers learn from experience. Both adaptive learning and punishment predict that equilibrium offers should be higher where second mover resistance to them is lower. Consistent with this, the rate at which first movers make equilibrium offers is higher in $? = 10$. But the difference is small (86% and 81%, respectively) and not statistically significant.

⁶ An experiment by Winter and Zamir (1997) finds that ultimatum second mover rejection rates are stable with repetition.

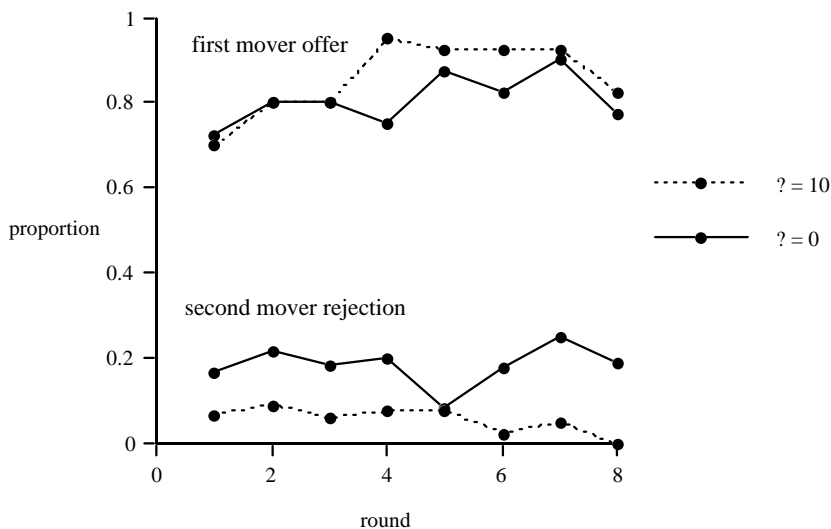


FIG. 4. Unequal split: Offers and rejections (each treatment: five sessions \times eight games/round).

Better evidence for learning comes from looking directly at the first mover adjustment pattern. Adaptive learning implies that first movers should be more likely to switch to an equal split offer after an unequal split is rejected than after it has been accepted. In fact, there is such a tendency. For each session, we took the rate at which an equal split was offered after an unequal split was accepted and subtracted this from the rate at which an equal split was offered after an unequal split was rejected (results ordered by session; no-rejection session 6 excluded):

0.28 0.59 0.08 -0.08 0.33 -0.02 -0.04 0.30 0.24.

All three negative entries are small, and all occur in sessions where there are relatively few rejections. A one-sample permutation test rejects the hypothesis of no switching tendency in favor of the hypothesis that switching is more likely after a rejected unequal split (one-tailed $p = 0.0234$).

More broadly, adaptive learning would lead us to expect that the probability of offering an unequal split increases after the unequal split is accepted, does not increase when it is rejected, and decreases after the equal split is accepted (the equal split was rejected just two times in the entire sample). To test this, we ran a probit, regressing first mover i 's offer

in round t (unequal split = 1) on i 's history of play through round $t - 1$; the history accounted for by the variables:

$R_{i,t-1}$ = number of times through round $t - 1$ that i offered the unequal split and it was rejected;

$A_{i,t-1}$ = number of times through round $t - 1$ that i offered the unequal split and it was accepted;

$E_{i,t-1}$ = number of times through round $t - 1$ that i offered the equal split (and it was accepted);

$\text{Treat}_i = 1$ if i was in a $? = 10$ and $= 0$ otherwise;

Random effects to account for heterogeneity among the first movers, i .

The estimated model (with two-tailed p -values in brackets) is

$$\hat{P}_t = -0.51R_{i,t-1} + 0.26A_{i,t-1} - 0.01E_{i,t-1} + 0.05\text{Treat}_i$$

$$\begin{array}{cccc} [0.003] & [< 0.001] & [0.916] & [0.862] \end{array}$$

$$\text{Log likelihood} = -218.2. \quad (3)$$

The coefficients for the history variables have the expected signs, although the coefficient for $E_{i,t-1}$ is not significant. We can benchmark the adaptive learning model of Eq. (3) against a simpler round learning model. To do so, first note that $R_{i,t-1} + A_{i,t-1} + E_{i,t-1} \equiv t - 1$ for all i . Then, if round learning is as good an explanation as (3), the history coefficients should not be significantly different from one another. The implied restriction is, however, easily rejected (Wald test, $p = 0.00001$).

Adaptive learning implies that the pattern of reinforcement, and therefore the pattern of learning, should be the same across treatments. (A difference in pattern might be taken as evidence for belief-based learning.) In fact, the coefficient for the Treat dummy is small and insignificant. We also tested for treatment differences in the coefficients for history variable by inserting the appropriate dummies; none were significant ($p > 0.400$ in all cases).

Summarizing the results in this section, we find that second mover rejection behavior differs across treatments as the punishment hypothesis would lead us to anticipate. That said, there are some rejections in the $? = 10$ treatment that the punishment hypothesis does not explain. One explanation is that in early rounds, $? = 10$ second movers reject to induce first movers to offer the equal split in later rounds, although the structure of the experiment makes it difficult to reconcile this explanation with conventional reputation-building models. There is no statistical evidence for second mover learning, although rejection rates in $? = 10$ drop in later

rounds. There is statistical evidence for first mover learning. The probability of offering the unequal split adjusts to the history of play in the manner adaptive learning predicts.

5. ADAPTIVE LEARNING: FITTING INITIAL PROPENSITIES FROM THE DATA

The second mover data from the experiment rejects the adaptive learning hypothesis as we developed it in Section 2. That development allotted no explanatory power to initial propensities (a nonlearning explanation). The difference in second mover rejection rates we observe is evident from the first round of play. It is natural to wonder whether a modified adaptive learning model, one that admits a difference in initial propensities, perhaps because the games have different fairness characteristics, would be consistent with the data. Adaptive learning might then provide a satisfactory account of the dynamics of repeated play.

In this section, we compare the simulated learning path from a simple adaptive learning model—fitting initial propensities from the data—to observed second mover behavior. The model’s predictions for first movers turn out to be sensitive to what we assume first movers believe about ?. Since there is no reliable way to independently gauge these beliefs, and since the results in the last section already demonstrate that first movers’ transitions between actions are at least roughly consistent with adaptive learning, we do not attempt any further fitting of first mover behavior. It turns out that second mover results are robust to what we assume about ?. The key test requires no assumption at all.

For the simulation, we use the 2×2 bimatrix version of the game displayed in Fig. 5. This differs from the complete normal form in that we have deleted the two second mover strategies in which the equal split is rejected. Since the equal split was rejected just two times in the entire sample, the deleted strategies are, empirically speaking, of negligible

		second mover	
		accept both	reject unequal
first mover	equal split	5 5	5 5
	unequal split	8 2	? 0

FIG. 5. Normal form game used for the simulations.

importance; deleting these strategies speeds the computer calculations considerably. The simplification is also equivalent to Roth and Erev's (1995) simplification for their simulations in the sense that both reduce the second mover strategy space to acceptance thresholds. In all other respects, we model the adaptive learning process much as described in Section 2.2.

5.1. *Average Learning Paths*

To get a feel for the simulation, we begin as is commonly done with these models, by looking at the average simulated learning paths over an extended time horizon (one far longer than the length of the experiment). The model requires us to fit two parameters: initial propensity and initial strength (see Section 2.2). For initial propensity, we use the rejection rate over the first two rounds of data, each treatment fitted separately. Since there is some difference of opinion over the proper way to fit initial strength—whether to use a value that has proven successful in previous experiments or a value that in some sense optimizes the model's fit with the data—we study the model over the plausible range of initial strength values. Because the value of β is unknown to the first mover, an assumption must be made about how it enters the propensities of his or her two strategies (that is, about the first mover's belief about the value of β). We suppose that whenever a first mover automaton observes a β , a fair coin is flipped to determine whether a payoff of 0 or 10 is attributed to the offer-an-unequal-split strategy. We have run simulations under a variety of schemes (e.g., "always 0," "always 10," or "always the expected value of 5"). With respect to second mover learning paths, the results are all similar.

Figure 6 shows the paths of average rejection rates of 2000 independent runs of the model simulation. Within each run, the original experimental matching scheme was iterated 500 times, and play was continued with current propensities across repetitions. After a brief rise at the beginning of the simulation, rejection rates slowly fall. Looking at the graph, the average paths of the two treatments appear rather similar, the major difference being initial play. It is unclear from the figure whether this is an adequate accounting of the differences we see in the actual data.

The simulation in Fig. 6 was generated using an initial strength of 10 (the same value used by Roth and Erev, 1995, p. 177). We have done average path simulations using initial strengths from 1.25 to 1280 with a grid factor of 2. Higher initial strengths slow learning and flatten paths. Lower initial strengths speed learning and make paths steeper. In no case, however, is there a clear difference between treatment paths.




FIG. 6. Average rejection rates: Adaptive learning model with initial strength set to 10. ($\theta = 0$ path begins at a frequency slightly above the peak frequency of the $\theta = 10$ path.)

All of the average path simulations (independent of initial strength) exhibit the sort of initial rise in rejection rates we observe in Fig. 6. The key factors here are the observed low initial propensity and the analysis' focus on averages. Specifically, given the low initial propensity to reject, if an equal split is offered to a second mover, this means that (1) it is attributed to the reject strategy with low probability, but (2) if it is attributed to it, the effect on the likelihood to play the reject strategy is far higher than if it is credited to the accept strategy. When the likelihood of the first mover offering an equal split is sufficiently high, the consequences of effect (2) are greater than those of effect (1), and the *average* propensity to reject the unequal split drifts upward. Simultaneously, however, first mover propensities to offer the unequal split rise, and at some point this induces the second mover average propensity to reject to reverse direction and fall. When we rerun the simulations using a higher initial propensity to reject, no drift is evident.

The observed paths are not directly comparable to the predictions in Fig. 6 because of the much longer time horizon of the simulations. We could restrict attention to the first eight rounds of the simulation—a real-time comparison. But such an analysis still ignores the large variance inherent to the model's predictions. The fact that the average simulated rejection path initially rises when the actual path is flat or falls (see Fig. 4), for instance, is not really a strike against the model: Even if the model's

behavioral assumptions are perfectly fulfilled, there is no reason to expect observed outcomes to exactly match the theoretical average path.

5.2. *The Corridor Test*


In this section, we apply a method of statistical comparison that is both real-time and accounts for theoretical variance. We take as null hypothesis that subjects behave according to the adaptive learning mechanism. We then ask, What second mover rejection rates are consistent with the null hypothesis in round t ? To compute this, we take the observed payoffs of each second mover up to round $t - 1$ as propensities for that subject's decision in round t . We estimate the distribution of possible rejection rates by simulation (10,000 repetitions). The result is a "corridor," a kind of confidence interval through which the model predicts second mover rejection behavior should move.

One advantage of the corridor test is that only *observed* first mover decisions are necessary to update second movers' propensities from round-to-round. First mover beliefs about the value of β play no role in the analysis. The procedure is somewhat biased in favor of the learning model since all past outcomes are interpreted as outcomes of an adaptive learning process, so systematic errors are corrected for, at least to some extent. This bias, however, serves to strengthen the conclusions we will draw.


Because the experiment was played in the extensive form, it cannot be detected whether the second mover has chosen the left or right column of the Fig. 5 matrix in response to an offer of an equal split. This problem, however, is easily solved since the model provides us with what are, by the null hypothesis, the correct underlying propensities. The resulting payoff is simply attributed randomly to one of the two strategies, with probabilities proportional to the current propensities. Second mover initial propensities are set as in the previous simulations.

Figures 7a and 7b show the rejection rates observed in the experiment along with the corresponding model predictions, derived from an initial strength of 10. The figures indicate the median of the predicted distribution and the bounds of its inner 90% interval. Under the null hypothesis, rejection rates below the lower line or above the upper line will not occur with a probability of more than 5%.

Visual examination shows that, for the relatively small sample sizes, the model predictions are not very sharp. A wide range of possible observations is compatible with the null hypothesis. But this aside, the model appears to work well for the $\beta = 0$ treatment, in which six of the eight observations are close to the median of the distribution. There is no



(a)



(b)

FIG. 7. Second mover rejections. Initial strength parameter set at 10.

systematic bias in prediction, and the null hypothesis cannot be rejected in any single round.

The picture is quite different for the $\tau = 10$ treatment, for which an overprediction of rejection rates can be observed. The data show a trend of declining rejection rates, but the model predicts modestly increasing frequencies. Since the model is initialized with data from the first two rounds, systematic biases can show up only in later rounds. In fact, the observed

rejection rates in all six rounds are below the median the model predicts. A sign test rejects the hypothesis of no systematic bias for the rounds the model predicts (two-tailed $p = 0.031$). In the last three rounds, observed rejection rates drop out of the 90% interval altogether.⁷

The initial strength for the corridor tests in Figure 7 was 10. It turns out that the results for $\gamma = 0$ are remarkably robust to the initial strength setting; the corridor and associated median are virtually unchanged for initial strengths from 0.001 to 100,000. For $\gamma = 10$, lowering the initial strength below 10 leads to faster learning, the predicted corridor and median rise over the first three rounds even more rapidly, and so the observed rejection path falls out of the corridor more quickly. In fact, we find that the $\gamma = 10$ data stays within the 90% corridor only if the initial strength is 100 or more. But even at very high initial strengths, the upward bias in the model's predictions, while reduced in size, persists. For example, at an initial strength of 1000, a sign test still rejects the hypothesis of no systematic bias for the last six rounds (two-tailed $p = 0.031$).

In sum, the model tracks rejection behavior in $\gamma = 0$ well but is insufficiently sensitive to the removal of the punishment opportunity in $\gamma = 10$ to capture the latter observed decrease in rejections. It appears that the punishment motive plays a role beyond initial propensities.

6. CONCLUSIONS

To summarize our findings, second movers were three times more likely to reject the unequal split when doing so punished the first mover than when doing so rewarded the first mover. There is no statistical evidence for second mover learning (although we observe the relatively low rejection rates in $\gamma = 10$ decrease further in later rounds). We do find statistical evidence for first mover learning. The probability of offering the unequal split adjusts to the history of play in the direction adaptive learning would predict. This raised the question of whether an adaptive model where initial propensities are fit conditional on the treatment to account for fairness effects might provide a comprehensive explanation of the data.

⁷ Second mover rejection rates in $\gamma = 10$ increase at first because the rate at which first movers offer the unequal split is not sufficient to penalize the initial propensity for second movers to reject the unequal split. The reader might think the rejection rates rise for the same reason we observed the initial upward drift in rejection rates in the average path simulations. But this is not quite correct. A quick way to see that something must be different is to note that a spike appears in both $\gamma = 0$ and $\gamma = 10$ average path simulations (Fig. 5), whereas early round rejections rise in the corridor test *only* in $\gamma = 10$ (Fig. 7a). The difference is that the corridor test uses *actual* propensities to generate each round's confidence interval, and this tends to dampen down the drift.

We found that, after adjusting for differences in initial propensities, a simple adaptive learning model has descriptive power for second movers in $\beta = 0$, the treatment most like the standard ultimatum contest. But the model overpredicts rejection rates of unequal splits in the $\beta = 10$ treatment.

By design, our experiment focuses its lens sharply on second movers. If we assume that second movers never reject a positive amount of money, then both theories would lead us to expect quick convergence to conventional perfect equilibrium play. In this sense, second mover behavior is the key to separating the two theories. At the same time, this focus also gives rise to the major limitations of our study. While we are able to identify some evidence that first movers learn in accord with the adaptive learning hypothesis, the design of the experiment makes it difficult to verify this by directly fitting the model to first mover data. Also, our experiment involves a modest number of iterations (eight rounds). While adaptive learning implies that learning should be more pronounced in earlier rounds, it is nevertheless possible that we might observe more learning over a longer time horizon.

Two main conclusions can be drawn from our results. First, a fairness motive, punishment, is a better explanation for why second movers reject some offers in the ultimatum game than is (slow) adaptive learning. Second, punishment appears to influence the dynamic path of the ultimatum game; simply fitting initial propensities to the adaptive learning model will probably be inadequate to account for this motive.

APPENDIX A: WRITTEN INSTRUCTIONS TO SUBJECTS (ORIGINAL TEXT IN GERMAN)

Player Types:

There are **two types** in the experiment: **player 1** and **player 2**.

After the introduction, each participant draws one of 16 cards.

The card drawn defines the terminal number of the participant.

The terminal number specifies the participant's type for **the whole experiment**.

Structure:

The experiment consists of **eight rounds**.

In each round there are eight **pairs of participants**, each with one *player 1* and one *player 2*.

In every round, every *player 1* meets a **different** *player 2*, and vice versa.

Thus, no **participant meets the same participant a second time**.

Decision:

In each round, in each pair, the following **two decisions** have to be made:

First, *player 1* chooses one of **two alternatives: Left or Right**.

Player 2 is informed about the choice of *player 1*.

Then, *player 2* also chooses one of **two alternatives: Left or Right**.

Payoffs in Talers:

<i>Player 1</i> chooses	<i>Player 2</i> chooses	<i>Player 1</i> receives	<i>Player 2</i> receives
Left	Left	5	5
Left	Right	0	0
Right	Left	8	2
Right	Right	?	0

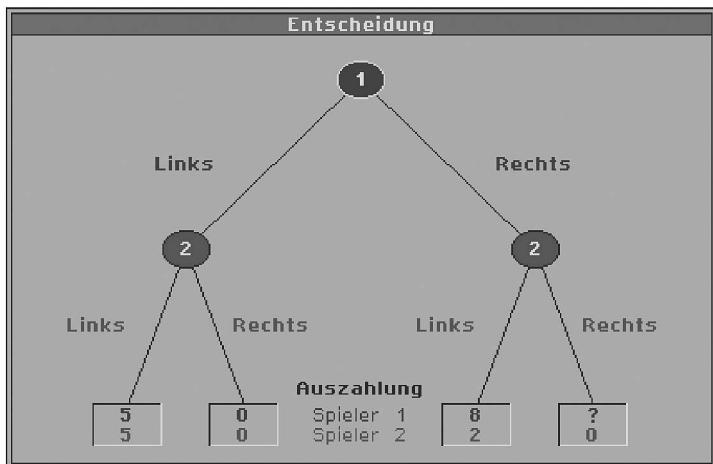
The Question Mark's Value:

The question mark **either has the value 0 or the value 10**, with each value being equally likely.

The question mark's value **does not change throughout the whole experiment**.

Only *player 2* knows the actual value of the question mark; *player 1* does not.

Player 2 are not allowed to announce this value.



Screen Shot of the Decision Tree

Exchange Rate:

Each taler is equivalent to **50 pfennigs**.

APPENDIX B: THE DATA

Each row represents a round with corresponding number. For “2R 16L” read “first mover, terminal 2, played R, and second mover, terminal 16, responded with L.” Sessions appear in the order they were run.

Session 1: ? = 10

1	2R 16L	11R 15R	13R 3L	8R 1L	14L 9L	17R 6L	12R 5L	18R 7L
2	2R 15R	11L 3R	13R 1L	8L 9L	14L 6L	17R 5L	12R 7L	18R 16L
3	2R 3L	11L 1L	13R 9L	8L 6L	14L 5L	17R 7L	12R 16L	18R 15L
4	2R 1L	11R 9L	13R 6R	8R 5L	14L 7L	17R 16L	12R 15L	18R 3L
5	2L 9R	11R 6L	13R 5L	8R 7L	14L 16L	17R 15R	12R 3R	18R 1L
6	2R 6L	11R 5L	13R 7L	8R 16L	14L 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16L	8R 15R	14L 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11R 16L	13R 15L	8L 3L	14L 1L	17R 9L	12R 6L	18R 5L

Session 2: ? = 0

1	2R 16L	11R 15L	13L 3L	8L 1L	14R 9L	17R 6L	12R 5L	18L 7L
2	2R 15L	11R 3L	13R 1R	8L 9L	14L 6L	17R 5R	12R 7R	18L 16L
3	2R 3L	11R 1L	13L 9L	8L 6L	14L 5L	17R 7R	12L 16L	18R 15L
4	2R 1L	11R 9L	13L 6L	8L 5L	14L 7L	17L 16L	12R 15L	18R 3L
5	2R 9L	11R 6L	13L 5L	8L 7L	14L 16L	17R 15L	12R 3L	18R 1L
6	2R 6L	11R 5R	13L 7L	8L 16L	14L 15L	17R 3L	12R 1L	18R 9L
7	2R 5R	11R 7R	13L 16L	8L 15L	14L 3L	17R 1L	12R 9R	18R 6L
8	2L 7L	11L 16L	13L 15L	8L 3L	14L 1L	17R 9R	12R 6L	18R 5R

Session 3: ? = 0

1	2L 16L	11R 15L	13L 3L	8L 1L	14R 9L	17R 6R	12R 5L	18R 7L
2	2R 15R	11R 3L	13R 1L	8R 9L	14R 6R	17R 5L	12R 7L	18R 16R
3	2L 3L	11R 1L	13R 9L	8R 6R	14R 5L	17R 7L	12R 16R	18R 15L
4	2R 1L	11R 9R	13R 6L	8R 5L	14R 7L	17R 16R	12R 15L	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7L	14R 16R	17R 15L	12R 3L	18R 1L
6	2R 6R	11R 5L	13R 7L	8R 16R	14R 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16R	8R 15L	14R 3L	17R 1L	12R 9L	18R 6R
8	2R 7L	11R 16R	13R 15L	8R 3L	14R 1L	17R 9L	12R 6R	18R 5L

Session 4: ? = 10

1	2R 16L	11L 15L	13R 3L	8R 1L	14R 9L	17R 6L	12R 5L	18R 7L
2	2R 15L	11R 3L	13R 1L	8R 9L	14R 6L	17L 5L	12R 7L	18R 16L
3	2R 3R	11L 1L	13R 9L	8R 6L	14R 5L	17R 7L	12R 16L	18R 15L
4	2R 1L	11R 9L	13R 6L	8R 5L	14R 7L	17R 16L	12R 15L	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7L	14R 16L	17R 15L	12R 3L	18R 1L
6	2R 6L	11R 5L	13R 7L	8R 16L	14R 15L	17L 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16L	8R 15L	14R 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11R 16L	13R 15L	8R 3L	14R 1L	17L 9L	12R 6L	18R 5L

Session 5: ? = 0

1	2R 16L	11R 15L	13R 3L	8R 1L	14L 9L	17R 6R	12L 5L	18R 7L
2	2R 15L	11R 3L	13R 1L	8R 9L	14L 6L	17L 5L	12R 7L	18R 16L
3	2R 3R	11R 1L	13R 9L	8R 6R	14L 5L	17L 7L	12R 16L	18R 15L
4	2R 1L	11L 9L	13R 6R	8L 5L	14L 7L	17L 16L	12R 15R	18R 3L
5	2R 9L	11R 6R	13L 5L	8R 7L	14L 16L	17R 15L	12R 3L	18R 1L
6	2R 6R	11R 5L	13L 7L	8R 16L	14L 15L	17R 3L	12R 1L	18R 9R
7	2R 5L	11R 7L	13R 16L	8R 15R	14L 3L	17R 1L	12R 9R	18R 6R
8	2R 7L	11R 16L	13R 15L	8L 3L	14L 1L	17R 9L	12R 6R	18R 5L

Session 6: ? = 10

1	2R 16L	11L 15L	13R 3L	8R 1L	14L 9L	17L 6L	12R 5L	18R 7L
2	2R 15L	11L 3L	13R 1L	8R 9L	14L 6L	17R 5L	12R 7L	18R 16L
3	2R 3L	11L 1L	13R 9L	8R 6L	14L 5L	17R 7L	12R 16L	18R 15L
4	2R 1L	11L 9L	13R 6L	8R 5L	14R 7L	17R 16L	12R 15L	18R 3L
5	2R 9L	11L 6L	13R 5L	8R 7L	14R 16L	17R 15L	12R 3L	18R 1L
6	2R 6L	11L 5L	13R 7L	8R 16L	14R 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11L 7L	13R 16L	8R 15L	14R 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11L 16L	13R 15L	8R 3L	14R 1L	17R 9L	12R 6L	18L 5L

Session 7: ? = 10

1	2R 16L	11L 15L	13L 3L	8L 1L	14R 9L	17L 6L	12R 5L	18L 7L
2	2R 15L	11L 3L	13R 1L	8R 9L	14R 6L	17R 5L	12R 7R	18R 16L
3	2R 3L	11R 1L	13R 9L	8R 6L	14R 5L	17R 7L	12R 16L	18R 15L
4	2R 1L	11R 9L	13R 6L	8R 5L	14R 7L	17R 16L	12R 15L	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7L	14R 16L	17R 15L	12R 3L	18R 1L
6	2R 6R	11R 5L	13R 7L	8R 16L	14R 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16L	8R 15L	14R 3L	17L 1L	12R 9L	18R 6L
8	2R 7L	11R 16L	13R 15L	8R 3L	14R 1L	17R 9L	12R 6L	18R 5L

Session 8: ? = 0

1	2R 16R	11R 15L	13L 3L	8R 1L	14R 9L	17R 6R	12R 5L	18L 7L
2	2R 15L	11R 3L	13L 1L	8R 9L	14R 6L	17R 5L	12R 7L	18R 16L
3	2R 3L	11R 1L	13R 9L	8R 6L	14R 5L	17R 7L	12R 16L	18R 15L
4	2R 1L	11R 9L	13R 6L	8R 5L	14R 7L	17R 16L	12R 15L	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7L	14R 16L	17R 15L	12R 3L	18R 1L
6	2R 6L	11R 5L	13L 7L	8R 16L	14R 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16L	8R 15L	14R 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11R 16L	13L 15L	8R 3L	14R 1L	17R 9L	12R 6L	18R 5L

Session 9: ? = 0

1	2R 16L	11R 15R	13R 3L	8R 1L	14R 9L	17L 6L	12R 5L	18R 7L
2	2R 15R	11L 3L	13R 1L	8R 9L	14R 6L	17L 5L	12R 7L	18R 16L
3	2L 3L	11R 1L	13R 9L	8R 6L	14R 5L	17R 7L	12R 16L	18R 15R
4	2R 1L	11R 9L	13R 6R	8R 5L	14L 7L	17L 16L	12R 15R	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7L	14R 16L	17R 15R	12R 3L	18R 1L
6	2R 6L	11R 5L	13R 7L	8R 16L	14R 15R	17L 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16L	8R 15R	14R 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11R 16L	13R 15R	8R 3L	14L 1L	17R 9L	12R 6L	18R 5L

Session 10: ? = 10

1	2R 16R	11R 15L	13L 3L	8L 1L	14R 9L	17R 6L	12R 5L	18R 7L
2	2L 15L	11R 3L	13R 1L	8R 9L	14R 6L	17R 5L	12R 7L	18R 16R
3	2L 3L	11R 1L	13R 9L	8L 6L	14R 5L	17R 7L	12R 16R	18R 15L
4	2R 1L	11R 9L	13R 6L	8R 5L	14R 7L	17R 16R	12R 15R	18R 3L
5	2R 9L	11R 6L	13R 5L	8R 7R	14R 16L	17R 15L	12R 3L	18R 1L
6	2R 6L	11R 5L	13R 7L	8R 16L	14R 15L	17R 3L	12R 1L	18R 9L
7	2R 5L	11R 7L	13R 16R	8R 15L	14R 3L	17R 1L	12R 9L	18R 6L
8	2R 7L	11L 16L	13L 15L	8R 3L	14R 1L	17R 9L	12R 6L	18R 5L

REFERENCES

- Abbink, K., and Sadrieh, A. (1995). "RatImage—Research Assistance Toolbox for Computer-Aided Human Behavior Experiments," University of Bonn, SFB Discussion Paper B-325.
- Blount, S. (1995). "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences," *Organ. Behav. Human Decision Processes* **63**, 131–144.
- Bolton, G. E. (1991). "A Comparative Model of Bargaining: Theory and Evidence," *Amer. Econ. Rev.* **81**, 1096–1136.
- Bolton, G. E., and Ockenfels, A. (2000). "ERC: A Theory of Equity, Reciprocity and Competition," *Amer. Econ. Rev.* **90**, 166–193.
- Bolton, G. E., and Zwick, R. (1995). "Anonymity versus Punishment in Ultimatum Bargaining," *Games Econ. Behav.* **10**, 95–121.
- Bush, R., and Mosteller, F. (1955). *Stochastic Models for Learning*. New York: Wiley.
- Camerer, C. F. (1997). "Progress in Behavioral Game Theory," *J. Econ. Perspec.* **11**, 167–188.
- Chen, Y., and Tang, F. (1998). "Learning and Incentive-compatible Mechanisms for Public Goods Provision: An Experimental Study," *J. Polit. Econ.* **106**, 633–662.
- Davis, D. D., and Holt, C. A. (1993). *Experimental Economics*. Princeton, NJ: Princeton Univ. Press.
- Duffy, J., and Feltovich, N. (1999). "Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study," *Int. J. Game Theory* **28**, 131–152.
- Erev, I., and Rapoport, A. (1998). "Coordination, Magic, and Reinforcement Learning in a Market Entry Game," *Games Econ. Behav.* **23**, 146–175.
- Erev, I., and Roth, A. E. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equilibria," *Amer. Econ. Rev.* 848–881.
- Fehr, E., and Schmidt, K. (1999). "A Theory of Fairness, Competition, and Cooperation," *Quart. J. Econ.* **114**, 817–868.
- Gale, J., Binmore, K. G., and Samuelson, L. (1995). "Learning to be Imperfect: The Ultimatum Game," *Games Econ. Behav.* **8**, 56–90.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). "An Experimental Analysis of Ultimatum Bargaining," *J. Econ. Behav. Organ.* **3**, 367–388.
- Harley, C. B. (1981). "Learning the Evolutionary Stable Strategy," *J. Theoret. Biol.* **89**, 611–633.
- Kagel, J. H., Kim, C., and Moser, D. (1996). "Ultimatum Games with Asymmetric Information and Asymmetric Payoffs," *Games Econ. Behav.* **13**, 100–110.
- Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *J. Econ. Theory* **27**, 245–252.
- Moir, R. (1998). "A Monte Carlo Analysis of the Fisher Randomization Technique: Reviving Randomization for Experimental Economists," *Experimen. Econ.* **1**, 87–100.
- Rabin, M. (1993). "Incorporating Fairness into Game Theory and Economics," *Amer. Econ. Rev.* **83**, 1281–1302.
- Roth, A. E. (1995). "Bargaining Experiments," in *Handbook of Experimental Economics* (J. Kagel and A. E. Roth, Eds.). Princeton, NJ: Princeton Univ. Press.

- Roth, A. E., and Erev, I. (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games Econ. Behav.* **8**, 164–212.
- Tang, F. (1996). "Anticipatory Learning in Two-Person Games: An Experimental Study. Part II: Learning," SFB Discussion Paper B-363, University of Bonn.
- Winter, E., and Zamir, S. (1997). "An Experiment with Ultimatum Bargaining in a Changing Environment," Working Paper 195, Washington University, School of Business and Center in Political Economy.