

Federal Reserve Bank
of Minneapolis

Spring 2000

Quarterly Review

The Suffolk Bank and the Panic of 1837 (p. 3)

Arthur J. Rolnick
Bruce D. Smith
Warren E. Weber

Learning to Be Unpredictable: An Experimental Study (p. 14)

Arijit Mukherji
David E. Runkle

Federal Reserve Bank of Minneapolis

Quarterly Review

Vol. 24, No. 2

ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

Any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Editor: Arthur J. Rolnick

Associate Editors: Patrick J. Kehoe, Warren E. Weber

Economic Advisory Board: Andrew Atkeson, V. V. Chari,

Patrick J. Kehoe, Ellen R. McGrattan

Managing Editor: Kathleen S. Rolfe

Article Editors: Kathleen S. Rolfe, Jenni C. Schoppers

Production Editor: Jenni C. Schoppers

Designer: Phil Swenson

Typesetter: Mary E. Anomalay

Production Assistant: Lucinda Gardner

Technical Assistant: R. Shawn Hewitt

Circulation Assistant: Elaine R. Reed

The *Quarterly Review* is published by the Research Department of the Federal Reserve Bank of Minneapolis. Subscriptions are available free of charge.

Quarterly Review articles that are reprints or revisions of papers published elsewhere may not be reprinted without the written permission of the original publisher. All other *Quarterly Review* articles may be reprinted without charge. If you reprint an article, please fully credit the source—the Minneapolis Federal Reserve Bank as well as the *Quarterly Review*—and include with the reprint a version of the standard Federal Reserve disclaimer (italicized above). Also, please send one copy of any publication that includes a reprint to the Minneapolis Fed Research Department.

Electronic files of *Quarterly Review* articles are available through the Minneapolis Fed's home page on the World Wide Web: <http://www.minneapolisfed.org>.

Comments and questions about the *Quarterly Review* may be sent to

Quarterly Review

Research Department

Federal Reserve Bank of Minneapolis

P. O. Box 291

Minneapolis, Minnesota 55480-0291

(Phone 612-204-6455 / Fax 612-204-5515).

Subscription requests may also be sent to the circulation assistant at elaine.reed@mpls.frb.org; editorial comments and questions, to the managing editor at ksr@res.mpls.frb.fed.us.

Learning to Be Unpredictable: An Experimental Study

Arijit Mukherji
Visitor
Research Department
Federal Reserve Bank of Minneapolis
and Associate Professor of Accounting
Carlson School of Management
University of Minnesota

David E. Runkle*
Research Officer
Research Department
Federal Reserve Bank of Minneapolis

Sometimes it pays to be unpredictable. Tennis players know this. So does the Internal Revenue Service. If a tennis player always served to the same part of the court, the player's opponent would know where to be in order to easily return the serve. And if the IRS published set rules about the factors that would trigger an audit, taxpayers would know how to easily avoid one. So the best strategy for both the tennis player and the IRS is to choose more or less randomly from among their alternative actions.

These kinds of strategic choices are often studied in noncooperative game theory. If all of the players in a game find it optimal to make their play unpredictable, the game is said to have a *mixed-strategy equilibrium*.

One standard assumption of noncooperative game theory is that all the players' payoffs are common knowledge. This is a fairly strong assumption. In many games, it is not true. But without this assumption, players of a game may not be able to play an equilibrium. Recent results in game theory show that players can reach equilibrium if they observe the actions of every player in the game and are able to use that information to infer the likely payoffs to other players, by using a particular learning process known as *sophisticated Bayesian learning* (Jordan 1991, Kalai and Lehrer 1993).

In this study, we use an experiment to test whether how much information players have available to them about the structure of a particular game affects their ability to play the mixed-strategy equilibrium of that game. We find that,

in our experiment, when players have all the information about the payoffs and actions of other players, they are able to play the game's mixed-strategy equilibrium. But players do not do that when they do not have complete payoff information, even when they have enough information about the game to possibly learn what they need to know about other players' payoffs. This result suggests that the players did not learn about the payoff structure of the game using sophisticated Bayesian learning. This result also suggests that economists should be careful about assuming in their models that people can easily infer everyone else's payoffs.

Games, Equilibrium, and Learning

Any game can be described by the actions of the players, the payoffs to those players, and the information that the players have about these actions and payoffs. Again, noncooperative game theory typically assumes *common knowledge*—that anything known about the game is known by all players, that each player knows what is known by all players, that each player knows that each player knows what is known by all players, ad infinitum. So, for example, if players in a game have complete information about the payoffs to all the players in the game,

*Also, Adjunct Professor of Finance, University of Minnesota.

common knowledge implies that each player knows that all players have this complete payoff information.

In game theory studies, the assumption of common knowledge is used to derive a Nash equilibrium, the principal solution concept in noncooperative game theory. A *Nash equilibrium* is a set of strategies for each of the players in the game, in which no player has an incentive to change his or her own strategy for choosing actions in the game as long as the other players do not change their strategies.

Unfortunately, game theory does not specify how players reach a Nash equilibrium or how they obtain the assumed common knowledge. If players have to learn about payoffs to other players, they might not arrive at the Nash equilibrium. Depending on the game, how players learn can have a great effect on whether they can play the Nash equilibrium.

Theories of the learning process in noncooperative games model the information players need to resolve strategic uncertainty in order to construct individually optimal strategies. One learning process that has been studied extensively is, again, sophisticated Bayesian learning.

Sophisticated Bayesian learning assumes that all players know their own payoff functions but not the payoff functions of other players. Uncertainty is represented as a subjective prior over other players' types, that is, as the determinants of other players' payoffs. Initially, strategies are chosen as a Bayesian Nash equilibrium of the static game. Then players observe the initial choices of all players, update their beliefs about the payoff functions of other players, play the Bayesian Nash equilibrium given the updated beliefs, and repeat the process. As play proceeds, all players update their beliefs about what determines the payoffs of the other players. As they attempt to figure out what characteristics of other players could be consistent with the optimality of those players' observed choices, players gradually learn the game, including the payoff functions of the other players.

Sophisticated Bayesian learning will always converge to Nash equilibrium, but it requires that players process a large amount of information.

In this study, we test whether subjects in an experimental game use sophisticated Bayesian learning to learn to play the game's Nash equilibrium. We manipulate the information available to the subjects through two experimental treatments. Then we test whether subjects' play is consistent with the predictions of sophisticated Bayesian learning.

A Game for the Experiment

For our test, we need to choose a game for which we can create an experimental treatment in which players can observe the actions of all other players without directly observing their payoffs. That is, players must have the minimum information necessary for sophisticated Bayesian learning. Two simple games are natural for an experimenter to consider for such a test: *two-person matching pennies* and *three-person matching pennies*. Only one of these games works for us.

Two-Person Matching Pennies

In a game of two-person matching pennies, each player can choose to play either *heads* or *tails*. If the chosen actions of the two players match, then player *A* wins player *B*'s penny, while if the actions do not match, then player *B* wins player *A*'s penny.¹ Table 1 shows how the actions of the two players are linked to their payoffs.

In this game, predictability does not pay. If player *A* always chose *heads*, then player *B* would always choose *tails* and win *A*'s penny. But if player *A* knew that player *B* would always choose *tails*, then player *A* would always choose *tails* and win *B*'s penny. Similarly, if player *B* knew that player *A* would always choose *tails*, then player *B* would always choose *heads*, to win player *A*'s penny. Each of the four combinations of actions gives one of the players an incentive to change actions. We thus say that this game has no equilibrium in *pure strategies*.

The game does have an equilibrium in another kind of strategy, however. Neither player in the two-person matching pennies game would be prepared to play *heads* or *tails* consistently, but each player would be willing to randomly play *heads* and *tails*, with equal probabilities for the two actions. If each player used this strategy independently of the other player, then the other player could not expect to earn more by changing actions.² The game has a *mixed-strategy equilibrium*.

Games with a mixed-strategy equilibrium are not just a theoretical nicety. Researchers have shown that mixed strategies are optimal in, for example, modeling financial and tax audits (Border and Sobel 1987 and Mookherjee and Png 1989) and monitoring work effort (Kanodia 1985). Thus, the two-person matching pennies game

¹This is a zero-sum game: the sum of the payoffs of player *A* and player *B* is zero.

²For a complete treatment of equilibrium in the two-person matching pennies game, see, for example, Gibbons 1992, pp. 33–34.

Table 1
How Actions and Payoffs Are Linked
in the Two-Person Matching Pennies Game . . .

Payoffs to Players A, B
When A Wins Matches and B Wins No Matches

		Player B's Action	
		Heads	Tails
Player A's Action	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

might appear to be a good one to use in an experiment on learning.

However, it isn't in ours. It doesn't work with sophisticated Bayesian learning. This type of learning requires that an experimental subject observe all the other subjects' actions but not directly observe their payoffs. But in a two-person game, if a subject observes the actions of the other subject and knows the payoff structure for the game, then the payoffs are obvious as well. For example, in the two-person matching pennies game shown in Table 1, if player A plays *heads*, B plays *tails*, A observes B's action, and A observes Table 1, then A knows that B's payoff was 1. Player A's knowledge of B's realized payoffs cannot be restricted without also restricting A's observation of B's actions or A's knowledge of the payoff structure of the game.

Three-Person Matching Pennies

The use of sophisticated Bayesian learning can be tested, however, with a three-person matching pennies game.

In a three-person game, all subjects' knowledge of other subjects' payoffs can be restricted while letting them see other subjects' actions. Consider the three-person matching-pennies game introduced by Jordan (1993) and shown in Chart 1. In this game, the three players, X, Y, and Z, can each choose either *heads* or *tails*. Each player has a counterpart: X's counterpart is Y, Y's counterpart is Z, and Z's counterpart is X. Each player's payoff is determined by whether the player's own choice of action matches the ac-

tion of his or her counterpart. If the player's action is the same as the counterpart's action, then the player receives nothing. If the player's action is different from the counterpart's action, then the player receives a penny.

Table 2 shows an example of actions and payoffs for the three players in this game. Suppose X chooses *heads*, Y chooses *heads*, and Z chooses *tails*. Because X's action matches Y's action, X receives nothing. Because Y's action does not match Z's action, Y receives a penny. Because Z's action does not match X's action, Z receives a penny.

With three players instead of two, the information that experimental subjects have about other subjects' payoffs can easily be varied so that subjects can observe the actions of all other subjects, even if they do not observe the payoffs to all subjects. Suppose that in a particular experimental treatment, subjects see the actions of all the other subjects. If subjects are told that every subject has a counterpart and subjects all know how matches are rewarded, then even if a subject observes the other subjects' actions, that subject will not know the other subjects' payoffs unless the subject sees Chart 1. For example, if X knows that X's counterpart is Y, but X does not observe Chart 1, then X does not know whether Y's counterpart is X or Z.

Not surprisingly, this three-person matching pennies game in Chart 1 has no equilibrium in pure strategies. If each player chose a pure strategy, one player would want to change actions.³ However, this game does have a unique Nash equilibrium in mixed strategies: each player chooses *heads* with probability $\frac{1}{2}$ independently. With this mixed strategy, the expected payoff to each player is half a penny. No player has an incentive to change the probability of choosing *heads* as long as each of the other two players also chooses *heads* with probability $\frac{1}{2}$.

Again we point out that sophisticated Bayesian learning converges to the Nash equilibrium. Therefore, if subjects use sophisticated Bayesian learning in repeated play of the three-person matching pennies game, then after enough periods, subjects should play the mixed-strategy Nash equilibrium.

³Suppose, as in Table 2, that X chooses *heads*, Y chooses *heads*, and Z chooses *tails*. Then Y and Z would each receive a penny, and X would receive nothing. Thus, X would have an incentive to change to the strategy of playing *tails*. Since X has this incentive to change, the assumed choice of actions cannot be an equilibrium. Now suppose that X changes his or her choice, but the others do not: X chooses *tails*, Y chooses *heads*, and Z chooses *tails*. Then X and Y would each receive a penny, and Z would receive nothing. Now Z would have an incentive to change from *tails* to *heads*, so this set of pure strategies cannot be an equilibrium. Similarly, it can be shown that no other set of pure strategies is an equilibrium for this game.

Chart 1 and Table 2

... And in the Three-Person Matching Pennies Game

Chart 1 Payoff Structure

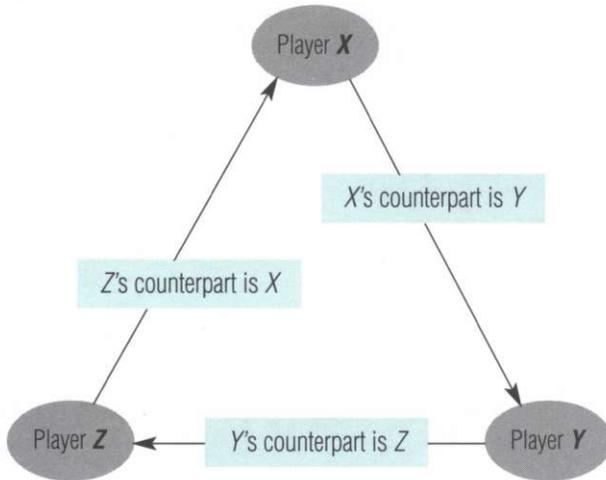


Table 2 Payoff Example

Payoffs When a Match Pays Nothing and No Match Pays a Penny

Player	Action	Payoff
X	Heads	0
Y	Heads	1
Z	Tails	1

The Experiment

Design

We designed an experiment to test whether controlling how much information that players have about the payoffs to other players affects their ability to play the mixed-strategy equilibrium of a version of the three-person matching pennies game.⁴

Our Game

The game in our experiment differs in four ways from the three-person matching pennies game described above.

- To make the payoff salient, we made the payoff for not matching the action of the counterpart a dime instead of a penny.⁵
- To avoid suggesting randomization, we changed the labels for the players' possible actions from *heads* or *tails* to the neutral *A* or *B*. (For simplicity, here, however, we will continue to use the labels *heads* and *tails*.)
- To minimize repeated-game effects, we put three players instead of just one at each node of the triangle in Chart 1, for a total of nine subjects.⁶ Each player thus had three counterparts, and the player's payoff was determined by the number of counterparts whose actions differed from the player's action. For example, if one person in group X chose *heads* while two people in group Y chose *tails* and one chose *heads*, then the person in group X failed to match two counterparts and received 20 cents. Depending on the actions of his or her counterparts, a player could earn 0, 10, 20, or 30 cents each time the game was played.
- As is common in economic experiments, we repeated the game for many periods with the same players, in order to see whether their behavior changed over time.⁷

Treatments

We designed two experimental treatments, which differ only by how much information players had about what determined the payoffs to other players.

In our first treatment, it was common knowledge that the subjects were matched according to Chart 1, which was shown and explained to the subjects in the common in-

⁴For a complete description and analysis of this experiment and its results, see McCabe, Mukherji, and Runkle 2000.

⁵With this payoff, subjects who played the equilibrium would earn at least \$12 per hour. This was more than twice the rate that the subjects—who were college students—could earn for unskilled on-campus employment at the time.

⁶This follows a suggestion of Friedman (1996).

⁷We did this to test whether subjects can learn to play the mixed-strategy equilibrium of the one-period game. In some games, the equilibria of the repeated game are not the same as the set of equilibria of the one-period game. Repetition can expand the set of strategies and payoffs; players can condition their behavior over time to punish or reward their counterparts. In this experiment, since there is a unique equilibrium in the one-period game, repeating it many times does not change the set of equilibria.

structions. The subjects also saw the entire history of choices for all subjects as it grew. The notion here was that if subjects' behavior in this baseline treatment is consistent with the mixed-strategy equilibrium, then we can legitimately investigate the role of payoff information by analyzing subjects' behavior in the other treatment.

In our second treatment, every subject knew from the common instructions that every subject's payoff depended on whether his or her choices matched those of three counterparts. However, nothing else was specified about those payoffs—who the counterparts were, for example. Over time, each subject observed the entire history of choices for all nine subjects. In this treatment, therefore, a subject had enough information to apply sophisticated Bayesian learning. While a subject did not know the incentives of the other subjects, these could be inferred from the history of play. If those inferences were made correctly, then the resulting observed behavior could be consistent with a mixed-strategy equilibrium. Sophisticated Bayesian learning is consistent with this possibility.

□ *Subjects, Sessions, and Periods*

As subjects for our experiment, we selected undergraduate students at the University of Minnesota. Each session in each treatment consisted of nine subjects who played the game for between 70 and 79 periods. We conducted several sessions for each of the two treatments, with 3,249 subject-actions in the first treatment and 3,312 subject-actions in the second treatment.⁸ For participating in the experiment, subjects were each paid \$5.00 plus earnings.⁹

In each session of the experiment, the subjects were separated by partitions, so that they could not see each other, and they were not allowed to talk to each other. After all the subjects had been seated, the experimenter displayed the instructions for the treatment on an overhead projector and read them to all subjects simultaneously. Then play began. Each session lasted between 60 and 90 minutes.

The experiment was run on a computer network. In each experimental period, the subjects saw on their computers information about the choices they and the others had made in the previous periods. (But, again, the amount of information the subjects had about the past varied by treatment.) In each period, subjects made their choices and entered them on their computers. After all subjects had made their choices in a particular period, their payoffs and the actions of all the other subjects were displayed.

Predictions

We tested whether controlling how much information subjects had about the structure of the matching pennies game affected subjects' play in the game. We tested two predictions: If the subjects in the experiment play a mixed-strategy equilibrium, then

- On average they should play *heads* with probability $\frac{1}{2}$.
- They should also randomize between *heads* and *tails* with probability $\frac{1}{2}$ regardless of the past actions of any of the subjects in the game. In particular, the probability that any subject chose *heads* in any given period should not depend on the number of *heads* chosen by the subject's counterparts in the preceding period.

For each of our two experimental treatments, we performed two tests to see whether the subjects played the mixed-strategy equilibrium of independently randomizing between *heads* and *tails* with a probability of $\frac{1}{2}$.

- We tested whether, on average, the subjects played *heads* with probability $\frac{1}{2}$.
- We tested whether, in a particular period, the probability that a subject chose *heads* was $\frac{1}{2}$, regardless of whether the subject's counterpart group had played zero, one, two, or three *heads* in the preceding period. For this test, we estimated four separate probabilities of playing *heads*—one for each possible number of *heads* played by the counterpart group in the preceding period—and tested whether the estimated probabilities were all $\frac{1}{2}$.

It seems reasonable to assume that subjects in the first experimental treatment, who have all the information about the structure of the game and can observe the actions and payoffs of all players in the game, will play the mixed-strategy Nash equilibrium. In fact, unless subjects play the mixed-strategy Nash equilibrium in the first treatment, the second treatment will not be interesting. Why test whether

⁸We randomized the number of periods so that the subjects would not know how many periods the experiments would last. This is consistent with O'Neill's (1987) experiment, in which subjects did not know how long it would last.

The average number of periods for the two experimental treatments was close to 73. If each of five sessions had had 73 observations, there would have been 3,285 ($5 \times 9 \times 73$) subject-actions in each experimental treatment.

⁹The subjects' median earnings were \$17.40; the 25th percentile of earnings was \$16.20, and the 75th percentile, \$18.00.

Charts 2–3

How Counterpart Play Affected a Subject's Actions

The Probability of a Subject Playing *Heads* in a Period vs.
The Number of *Heads* Played by Counterparts in the Preceding Period

Chart 2 With Complete Payoff Information

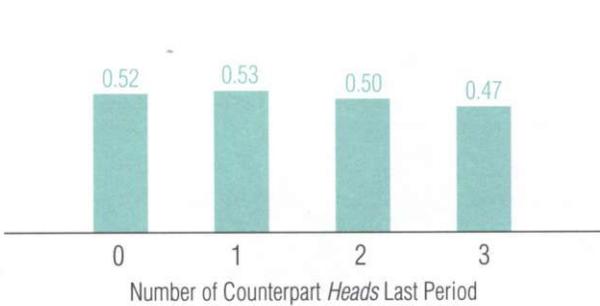
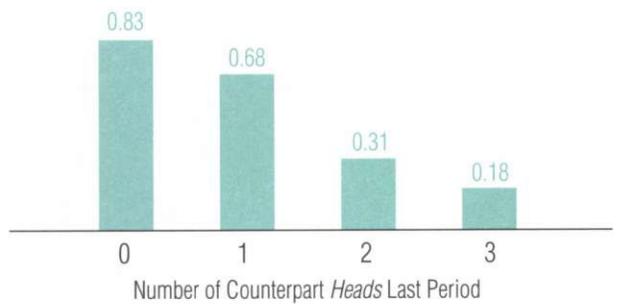


Chart 3 With Incomplete Payoff Information



subjects with limited information play the mixed-strategy Nash equilibrium if subjects with complete information do not?

Subjects in the second experimental treatment have limited information about the payoff structure of the game, but they observe the actions of all subjects. Therefore, they could use sophisticated Bayesian learning to determine the best strategy for their play. If all subjects used sophisticated Bayesian learning in this treatment, then they would randomize, that is, play the mixed-strategy Nash equilibrium.

Results

In our experiment, subjects do not seem to have used sophisticated Bayesian learning.

The unconditional probability that the subjects played *heads* was about the same in the two experimental treatments. In both, subjects came close to unconditionally randomizing between *heads* and *tails* with equal probability. (The probabilities are 0.508152 for the first treatment and 0.500468 for the second.) This result should not be surprising, at least for the first treatment. Subjects in the first treatment have complete information, so we expect them to play the mixed-strategy Nash equilibrium and randomize between *heads* and *tails* with equal probability. If sub-

jects in the second treatment used sophisticated Bayesian learning, then overall they would also randomize between *heads* and *tails* with equal probability.

However, the probability of playing *heads* in the two treatments was quite different when subjects' actions were conditioned on the previous actions of their counterparts.

In the first experimental treatment, remember, the subjects saw all the actions of the other subjects and had all the information they needed to compute the payoffs to all those other subjects. Here the subjects played *heads* and *tails* with nearly equal probabilities, regardless of the number of *heads* played by their counterpart groups in the preceding period. Chart 2 shows that in this treatment, the probability that a subject played *heads* was about 0.52 if no counterpart had played *heads* in the preceding period and about 0.47 if all three counterparts had played *heads* in that period. Thus, subjects in this treatment do seem to have chosen *heads* and *tails* with equal probability regardless of counterpart play in the preceding period. We tested statistically the hypothesis that the true probability of choosing *heads* in this treatment was $\frac{1}{2}$, regardless of previous counterpart choices. The test failed to reject this hypothesis at the 5 percent level of significance.

Subjects acted differently in the second experimental treatment. In this treatment, remember, subjects were al-

lowed to see all other subjects' actions, but they were not told the rules necessary to determine all the other subjects' payoffs.¹⁰ Chart 3 shows that in this treatment, the probability that a subject played *heads* was about 0.83 if no counterpart had played *heads* in the preceding period, but was only about 0.18 if all three counterparts had played *heads*. Thus, subjects in this treatment do not seem to have chosen *heads* and *tails* with equal probability regardless of counterpart play in the preceding period. We statistically tested the hypothesis that the true probability of choosing *heads* in this treatment was $\frac{1}{2}$, regardless of previous counterpart choices. The test strongly rejected that hypothesis.

If the subjects in the second treatment had used sophisticated Bayesian learning to determine what their best strategy should be, then they would have played the mixed-strategy Nash equilibrium of choosing *heads* with probability $\frac{1}{2}$ regardless of previous counterpart choices. The fact that the subjects did not randomize actions with equal probability, regardless of previous counterpart choices, shows that they did not use sophisticated Bayesian learning in this experiment.

Conclusions

In our experiment, information affects players' ability to play a game's mixed-strategy equilibrium. The only difference between our two experimental treatments is how much information the subjects have about the payoff structure for all the other subjects. In the first treatment, subjects were told all about that payoff structure, while in the second treatment, they were not. Even though this difference between the treatments should have had no effect on the subjects' ability to randomize if subjects had used sophisticated Bayesian learning, the subjects' play in the two treatments was significantly different. With complete payoff information, subjects did randomize and play the mixed-strategy equilibrium. Without complete payoff information, they did not. We conclude, therefore, that subjects in the second treatment did not use sophisticated Bayesian learning to find their best strategy, which was to randomize between the two actions with equal probability.

These results confirm the importance of information for playing mixed strategies. The results also suggest that economists should rethink models that assume people know the payoffs to everyone in the model. Since that is seldom true in real life, people often must infer payoffs after observing other people's actions. However, as this study demonstrates, people may not make those inferences correctly, even when theory suggests that they could.

References

- Border, Kim C., and Sobel, Joel. 1987. Samurai accountant: A theory of auditing and plunder. *Review of Economic Studies* 54 (October): 525–40.
- Friedman, Daniel. 1996. Equilibrium in evolutionary games: Some experimental results. *Economic Journal* 106 (January): 1–25.
- Gibbons, Robert. 1992. *Game theory for applied economists*. Princeton, N. J.: Princeton University Press.
- Jordan, J. S. 1991. Bayesian learning in normal form games. *Games and Economic Behavior* 3 (February): 60–81.
- . 1993. Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior* 5 (July): 368–86.
- Kalai, Ehud, and Lehrer, Ehud. 1993. Rational learning leads to Nash equilibrium. *Econometrica* 61 (September): 1019–45.
- Kanodia, Chandra S. 1985. Stochastic monitoring and moral hazard. *Journal of Accounting Research* 23 (Spring): 175–93.
- McCabe, Kevin A.; Mukherji, Arijit; and Runkle, David E. 2000. An experimental study of information and mixed-strategy play in the three-person matching-pennies game. *Economic Theory* 15 (no. 2): 421–62.
- Mookherjee, Dilip, and Png, Ivan. 1989. Optimal auditing, insurance, and redistribution. *Quarterly Journal of Economics* 104 (May): 399–415.
- O'Neill, Barry. 1987. Nonmetric test of the minimax theory of two-person zero-sum games. *Proceedings of the National Academy of Sciences of the United States of America* 84: 2106–9.

¹⁰In this experimental treatment, recall, it is common knowledge that each subject is matched with three other subjects, but subjects did not know how they were matched. There are many ways in which nine subjects can be matched with one another while still matching each subject with three other subjects. Our research tested whether subjects learned the actual structure of the payoffs (shown in Chart 1) by observing the past actions of other subjects. Jordan (1991) shows that if subjects follow a particular learning rule, then they will eventually converge to an equilibrium. In our case, this would be the unique mixed-strategy equilibrium.