

Decentralized Organizational Learning: An Experimental Investigation

By ANDREAS BLUME, JOHN DUFFY, AND APRIL M. FRANCO*

We experimentally study decentralized organizational learning. Our objective is to understand how learning members of an organization cope with the confounding effects of the simultaneous learning of others. We test the predictions of a stylized, rational agent model of organizational learning that provides sharp predictions as to how learning members of an organization might cope with the simultaneous learning of others as a function of fundamental variables, e.g., firm size and the discount factor. While the problem of learning while others are learning is quite difficult, we find support for the comparative static predictions of the model's unique symmetric equilibrium. (JEL C72, D23, D83)

If one's own actions are embedded in an ecology of the actions of many others (who are simultaneously learning and changing), it is not easy to understand what is going on. The relationship between the actions of individuals in the organization and overall organizational performance is confounded by simultaneous learning of other actors.

— Daniel A. Levinthal and James G. March (1993, 97)

In this paper we experimentally study *decentralized organizational learning*. We look at the team, the smallest organizational unit, and formulate a team-learning problem that is *decentralized* in that members of the team make their decisions independently. Team members have common interests and their objective is to discover through trial and error which of the combinations of team members' actions yield a positive payoff, without being able either to communicate with other team members or to observe their actions.

We are interested in how learning members of a team cope with the confounding effects of the simultaneous learning of other team members.¹ This issue is a recurrent theme in the literature on organizational learning, e.g., Pertti H. Lounamaa and James G. March (1987) and Daniel A. Levinthal and March (1993). It arises because, unless there is either unrestricted communication

* Blume: Department of Economics, University of Pittsburgh, 230 S. Bouquet St., Pittsburgh, PA 15260 (e-mail: ablume@pitt.edu); Duffy: Department of Economics, University of Pittsburgh, 230 S. Bouquet St., Pittsburgh, PA 15260 (e-mail: jduffy@pitt.edu); Franco: Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S 3E6 (e-mail: April.Franco@Rotman.Utoronto.Ca). Funding for this project was provided by the National Science Foundation. We thank the two referees, along with David Laibson, Matt Mitchell, and Avi Goldfarb, for helpful comments and suggestions. We have also benefited from presenting this paper at Carnegie-Mellon University and Harvard University, and at meetings of the Econometric Society, the Economic Science Association, and the Society for Economic Dynamics. Samvel Atayan, Adam D'Angelo, Charles Kannair, and Scott Kinross provided expert research assistance.

¹ An organization is any collection of individuals linked by a common purpose (Douglass C. North 1994). This includes naturalistic organizations (like firms and plants), but also teams and groups (within a firm). In our experimental investigation, we focus on groups of agents in the laboratory. As Linda Argote (1999, 99) notes, "understanding how groups or teams learn to work effectively together provides the micro foundation for organizational learning because groups or teams are the building blocks of most organizations."

or perfect observability, it can be difficult to attribute success or failure to combinations of actions taken by different members of the organization. For example, innovations in a firm's research and development department may be falsely regarded as ineffective, only because they are not effectively communicated to, and matched by changes in, the marketing department. Indeed, Clayton M. Christensen (1997) documents how a lack of communication between engineers and marketing departments at companies manufacturing computer disk drives led to the failure of some of the more established manufacturers. Similar failures to communicate, e.g., between nurses and physicians or within nursing teams, are cited by Amy C. Edmondson (2004) for the frequent lack of learning from failure in health care organizations. One of the key findings from Edmondson's empirical study is that "process failures in hospitals have systemic causes, often originating in different groups or departments from where the failure is experienced, and so learning from them requires cross-departmental communication and collaboration." Our experimental team-learning task captures this lack of communication and observability as starkly as possible. Furthermore, we compare multiplayer team learning with individual learning, where communication and observability constraints are removed.

Much of the work on organizational learning follows Herbert A. Simon's (1947) and Richard M. Cyert and March's (1963) views that the organization is composed of boundedly rational agents. The question then is how such organizations perform as a function of certain exogenously specified individual learning rules. For example, Richard H. Day and E. Herbert Tinney (1968) study decentralized learning in a firm with two independent decision makers who respond to success or failure by modifying their decision rules until a satisficing criterion is met. Day and Tinney examine the roles of "caution," "daring," and "failure response" and show that "learning—tempered by caution in response to failure" can eventually solve the firm's problem. Outside of economics, computer scientists have shown considerable interest in boundedly rational concurrent learning in multiagent environments.²

Our approach toward obtaining theoretical predictions, based on Blume and Franco (2007), is novel in the organizational learning literature, as we consider fully rational agents and use an explicit game theoretic framework. Rather than exogenously specifying individual learning rules, we predict learning behavior and tie our predictions to parameters governing individual preferences as well as to properties of organizations. Thus, by contrast with the bounded rationality approach where the focus is on *learning an equilibrium*, we investigate *learning in equilibrium*. Using the terminology of Day and Tinney, we can predict agents' "caution," "daring," and "failure response" from their equilibrium values. This novel approach complements the traditional approach to decentralized learning in organization theory and multiagent learning in computer science of exogenously specifying learning rules for individual agents. Our model yields sharp, testable predictions about behavior in the organization and about how this behavior varies with the fundamental variables that characterize the organization.³

Our focus is on a benchmark environment characterized by decentralized learning. In this environment, there is no room for explicit coordination, and due to limited information feedback, there is little room for tacit coordination. This implements some of the important hurdles faced by Jacob Marschak's (1960) "several-person firm," where each decision maker "decides about different things and on the basis of different information." Organizations have a multitude

² For example, Sandip Sen and Mahendra Sekaran (1998) investigate reinforcement learning in multiagent systems.

³ Fully rational learning in strategic settings is also investigated in the literature on many-agent versions of the multiarmed bandit problem (Patrick Bolton and Christopher Harris 1999), games with unknown payoff distributions (Thomas Wiseman 2005), and the literature on informational herding (e.g., Abhijit V. Banerjee 1992; Sushil Bikhchandani, David Hirshleifer, and Ivo Welch 1992). Our environment differs from those examined in these papers due to our focus on how agents with common interests coordinate their learning activities when there are constraints on their ability to communicate and to observe each other's actions.

of coping strategies for dealing with the problems arising from decentralized decision making (and dispersed information). Nevertheless, some of the decision making in an organization will remain decentralized. Also, we are likely to better understand the coping strategies and their value by first studying the extreme of no explicit coordination.

In the theoretical analysis for our environment, decentralization is captured through explicit constraints on agents' equilibrium strategies. We base our predictions on those equilibria that make no initial role distinctions among agents, on the grounds that such role distinctions are unlikely to be achieved without some explicit coordination mechanism. This approach, which models absence of a common language that could be used to distinguish roles by requiring that strategies respect the symmetries of players and actions, was pioneered by Vincent P. Crawford and Hans Haller (1990) and further developed by Blume (2000), Francis Kramarz (1996), V. Bhaskar (2000) and Steve Alpern and Diane J. Reyniers (2000).⁴ In future work we intend to use this benchmark environment as a platform for investigating routines, communication, information systems, culture, etc., in organizations. For example, in our environment, there are efficient routines that rely on role distinctions among agents.

Our principal experimental finding is that in their own learning behavior agents appear to take into account the confounding effects of the simultaneous learning of others. The data suggest that agents' learning behavior is sensitive to both group size and induced time preference. The direction of these comparative statics effects, although not their magnitude, is as predicted by theory.

The paper is organized as follows. In Section I we briefly review the theory that motivates our experiment and describe the behavior predicted by theory. In Section II we describe our experimental design. Section III summarizes our predictions and Section IV reports the aggregate experimental results. In Section V we examine the extent to which three different behavioral models might explain our aggregate findings, and in Section VI we briefly explore individual behavior. Conclusions are offered in Section VII.

I. A Model of Decentralized Learning

Our experimental treatments implement a class of learning situations analyzed in Blume and Franco (2007). They consider a game in which a collection of agents jointly try to find an optimal action combination in the face of limited information feedback, without being able to communicate and without reliance on a priori role distinctions. In this section we describe this game and the solution of the game that gives us our experimental predictions.

In the *search-for-success* game, n players repeatedly play a stage game in which each player has an identical number, m , of actions. All action combinations are either *failures* or *successes*. The (normalized) payoff for each player in the organization from a failure is zero and from a success is one. There are k success profiles and the remaining profiles are failures. Each assignment of the k successes to profiles is equally likely. Players know k but not the assignment of successes to profiles.

In the repeated game the random assignment of successes to action profiles is determined once and for all before the first play of the game. The stage game is repeated in rounds $t = 1, \dots, T$, until either a success is played once or the time horizon T is reached. Players observe only their own actions and their own payoffs, not the actions of the other players. Players maximize the expected present discounted value of future payoffs with discount factor δ , where $\delta > 0$. In the

⁴ In the experimental literature, symmetry constraints of this type have been studied by Blume et al. (1998, 2001), Blume and Uri Gneezy (2000, 2001), and Roberto A. Weber and Colin F. Camerer (2003). Weber and Camerer use a language construction experiment to study conflicting organizational cultures in the laboratory.

experiment, it is convenient to permit values of $\delta > 1$ in order to generate salient payoff differences; this can be interpreted as the learning task becoming more productive over time.

Decentralization is captured through a symmetry constraint on agents' joint strategies. In our setting, players are a priori identical, i.e., they have identical action sets, information, and payoff functions. Furthermore, there is no preplay communication or alternative mechanism that could help to desymmetrize players. Since there is nothing that distinguishes players a priori, we believe it is natural to expect that players hold identical beliefs about each other at the beginning of the games. Specifically, we study optimal symmetric strategies and symmetric equilibrium strategies in the repeated game. This approach was pioneered by Crawford and Haller (1990).

In the general framework, one can show that for any set of parameters and any length of the game, optimal symmetric strategies exist and optimal symmetric strategies are Nash equilibria. One can also show that optimal symmetric strategies are complex. There does not exist a symmetric public Nash equilibrium. In the infinite horizon game, any length of time before a success is found has positive probability. Optimal symmetric strategies can never be either completely deterministic or completely random. A central property of a solution is that agents invest in desymmetrization. In particular, they sacrifice current payoffs in order to increase the likelihood of reaching an asymmetric history: even before players have exhausted all of their actions, they will, with positive probability, return to action profiles they have visited before.

For the purpose of the experiment, we will focus on three-round versions of the game in which each agent has $m = 2$ actions. In that case, one can show that the game has a unique symmetric Nash equilibrium. As in the general case, in this equilibrium both random switching and deterministic switching are part of the equilibrium strategy. Blume and Franco (2007) show that the unique symmetric Nash equilibrium has the following form. In the first round agents randomize uniformly over their two actions. In the second round they switch to a different action with probability

$$p^*(\delta, n, k) = \frac{2}{\left(\frac{1}{\left(\frac{1-\delta}{\delta} \right) \left(\frac{2^n - 2}{2^n - 1 - k} \right) + 1 \right)^{1/(n-1)} + 2},$$

which depends on the discount factor δ , the number of agents n , and the number of success profiles, k . In the third round they switch with probability $q_0 = 1$, if they didn't switch in the previous round, and with probability $q_1 = 0.5$ if they did switch in the previous round.

The intuition underlying this solution is easily understood in the two-player, two-action case. Refer to the two players as player A and player B (only for record keeping purposes, without implying any asymmetry among the players). In the third round, conditional on not having switched before, it is clearly optimal for player A to switch with probability $q_0 = 1$ since this is the only way to be entirely sure that a new action combination will be examined. Conditional on having switched before, player A faces two possibilities. The first possibility is that player B has not switched. Then, by the argument just given, player B will switch in the third round with probability one, guaranteeing that a new action combination will be examined, irrespective of player A's choice. The other possibility is that player B has switched. In that case, if both players use identical switching probabilities q_1 , it is easily seen that the value of this switching probability that maximizes the probability of a novel action profile is $q_1 = 0.5$. The remaining question is why agents randomize with nondegenerate probabilities in the second round. Observe that there cannot be a symmetric equilibrium in which both agents switch with probability one in round 2. Otherwise player A could simply stay put in round 2 and then switch with probability one in round 3. This deviation would guarantee that in each round a novel action profile would be visited. Therefore the deviation would be profitable, breaking the putative equilibrium. Since in

this game an optimal symmetric strategy has to be an equilibrium, this argument also shows that probability-one switching in the second round cannot be optimal (provided that $\delta > 0$).

From the explicit formula for $p^*(\delta, n, k)$ we can derive a number of interesting testable comparative statics predictions: $p^*(\delta, n, k)$ is (i) strictly decreasing in δ , (ii) strictly decreasing in n , (iii) strictly decreasing in n even if $k/2^n$ is kept constant, and (iv) strictly increasing in k .

II. Experimental Design

In our experiment participants repeatedly play a search-for-success game. Each experimental session involved 16 *periods* of play. In each period, subjects played a *three-round* search-for-success game. Before the beginning of any period, participants are randomly (re)matched into groups of fixed size. During a period, all participants belonging to the same group played the three-round search-for-success game with one another.

We employ a 3×2 experimental design. The first treatment variable is the number of individuals who participate as a team to play the search-for-success game, either one individual (“singles”), two individuals (“pairs”), or three individuals (“triples”). This allows us to investigate the distinction between individual and organizational learning and whether individual learning within an organization is sensitive to the size of that organization. For each of these three treatments, we consider a second treatment variable, the order in which we vary the discount factor: eight periods with a high discount factor followed by eight periods with a low discount factor (*high-low*), or the reverse order (*low-high*). The precise details of our parameterization of the model and experimental design are provided in Table 1.

Representatively, for the six treatments, consider the details of the *pairs-low-high* treatment. In each of the two sessions of this treatment, a cohort of 20 students was recruited from the undergraduate population of the University of Pittsburgh. None of these students had prior experience with any of the treatments in this experiment. The students were randomly assigned to separate computer terminals and received written instructions. The instructions for the experiment were read aloud to the students to make them common information.⁵ Prior to the start of period 1, subjects were randomly matched into ten pairs. During period 1, each pair played a three-round search-for-success game in which each individual chose between two actions, X and Y . Subjects were informed that their initial choice, say (X, Y) , was payoff-irrelevant and served only to determine the remaining three profiles, here (X, X) , (Y, X) , (Y, Y) , each of which was then equally likely to be the unique success profile, $k = 1$.⁶ Players were informed of this procedure for selecting the success outcome, but were not told which of the three outcomes had been chosen as the success outcome. Consistent with the theory, players were not given any information about the choice of their match in the first round or in any subsequent round of the game. They did, of course, know the action they chose in each round. In the second round, participants were prompted via their computer terminal to enter the probability with which they would switch from their first-round choice of X or Y to the other choice. Players could specify any probability in $[0, 1]$ (up to six digits) representing the probability with which they would like to switch. Players were informed that a choice of zero insured that no switch would occur, while a choice of one insured that a switch in action would definitely occur. A choice in $(0, 1)$ meant that a switch would occur if a random number drawn by the computer program was less than or equal to the player’s chosen

⁵ The instructions used in this treatment and the other five treatments (along with other supporting materials) are provided in the web Appendix (available at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.4.1178>).

⁶ The reason we made the first-round payoff irrelevant is that we wanted to maximize the number of observations we obtained on individual decisions regarding action choices in the search-for-success game. While our decision to make the first-round payoff irrelevant should not affect the second- and third-round behavior of rational agents, it remains an open question whether this is in fact the case.

TABLE 1—CHARACTERISTICS OF EXPERIMENTAL SESSIONS

Design choice	Team size treatments		
	Singles	Pairs	Triples
Number of players on a team (n)	1	2	3
Number of action choices (m)	4	2	2
Number of success profiles (k)	1	1	2
Number of rounds per game (T)	3	3	3
Number of sessions with low-high order: $\delta = 0.5$, $\delta = 6^a$	2	2	2
Number of sessions with high-low order: $\delta = 6$, $\delta = 0.5^b$	2	2	2
Number of subjects per session	10	20	15
Number of teams per session	10	10	5
Total number of subjects	40	80	60
Total number of teams	40	40	20

^aEach session consists of 8 three-round periods with $\delta = 0.5$ followed by 8 three-round periods with $\delta = 6$.

^bEach session consists of 8 three-round periods with $\delta = 6$ followed by 8 three-round periods with $\delta = 0.5$.

cutoff value; otherwise, no switch would occur.⁷ After all players had specified their round 2 switching probabilities, the computer program determined whether each player switched or not. Each player was then informed of their *own* action choice for round 2 and whether they and their partner had achieved the success outcome in round 2. If the success profile was chosen, both participants in that pair received a payoff of \$1, which ended their search-for-success game for that period. Otherwise they proceeded to round 3 and were prompted, via their computer monitor, to enter the probability with which they switched their action from their *second*-round choice to the other choice. Following submission of this probability, they were informed only of their own action choice for round 3. If a pair of players' choices resulted in finding the success profile in the third round, both received a payoff of $\delta\$1 = \0.50 as we chose to set $\delta = 0.5$ in this treatment. Otherwise both received a payoff of \$0. Following the third round of the first period, that period was declared over; all participants were randomly rematched into pairs and proceeded to play the same three-round search-for-success game in period 2 in their new pairings. The computer program ensured that a player's match in the current three-round game (period) differed from his match in the previous three-round game (period), and this fact was made known to subjects. This pattern was repeated for eight periods. Following completion of the eighth period, subjects were instructed that for the remaining eight periods (three-round games) the discount factor would be raised to $\delta = 6$ so that achievement of the success profile in round 2 continued to pay \$1 while a third-round success now paid \$6. The change in δ was not announced until the start of the ninth period. The new discount factor was in effect for the remaining eight periods of the session (periods 9–16). At the end of the session, subjects were paid their earnings from all 16 periods (three-round games) played in cash in addition to a guaranteed \$5 show-up fee.

The differences in the remaining five treatments were as follows (see also Table 1). In the *high-low* treatments, the discount factor was 6 in the first eight periods and 0.5 in the second eight periods. In the *singles* treatments ($n = 1$), each individual plays alone, and instead of choosing between actions X and Y he chooses among the profiles (X, X) , (X, Y) , (Y, X) , (Y, Y) , that is, $m = 4$ rather than 2. As in the pairs treatment there is a single success outcome ($k = 1$) in the singles treatment. In the *triples* treatments ($n = 3$), members of each three-player group choose between X and Y ($m = 2$) and, instead of one success profile, there are two, $k = 2$; the latter choice keeps the ratio of k to m^n constant across treatments.

⁷ While it was not necessary to provide subjects with a randomization device, and it may have encouraged subjects to randomize more than they otherwise would have chosen to, we nevertheless thought the benefits of providing such a device, given that the symmetric equilibrium typically involves playing a mixed strategy, outweighed any costs.

The average payoff across all treatments and all sessions earned by subjects, including the \$5 show up fee, was \$25.37 for an experimental session that lasted approximately 75 minutes.

III. Predictions

In this section we describe the point predictions of the theory for our experimental treatments and use these to formulate hypotheses regarding the comparative statics effects of changing the number of players n and the discount factor δ , and about the relative magnitude of conditional and unconditional switching probabilities.

For the *singles* treatment, the formula from Section II does not apply, but the optimal solution is easily derived. If an individual has not achieved a success in round 2 it is clearly optimal to switch with probability one in round 3, regardless of whether there was a switch in round 2 or not, i.e., $q_0 = q_1 = 1$. Given these conditional switching probabilities, the expected payoff from switching in round 2 equals $1/3 + 2/3 \delta 1/2$, where $1/3$ is the probability of a success in round 2, and $1/2$ is the probability of a success in round 3 conditional on not having succeeded in round 2. In contrast, the expected payoff from not switching in round 2 equals $0 + 1/3 \delta$, which shows that with $n = 1$, switching is preferable to not switching regardless of the discount factor. Therefore, we have: for $n = 1$, $k = 1$, and any positive δ , theory predicts the vector of switching probabilities $(p, q_0, q_1) = (1, 1, 1)$.

For the *pairs* treatment, the unconditional switching probability p can be calculated as a function of δ from the formula in Section II. Recall the conditional switching probabilities $q_0 = 1$ and $q_1 = 0.5$ are in this case independent of δ . This gives us: for $n = 2$, $k = 1$, and $\delta = 0.5$, theory predicts the vector of switching probabilities $(p, q_0, q_1) = (0.8, 1, 0.5)$; for $n = 2$, $k = 1$, and $\delta = 6$, theory predicts the vector of switching probabilities $(p, q_0, q_1) = (0.25, 1, 0.5)$.

For the *triples* treatment, we can once again calculate p from the formula in Section II, and the conditional switching probabilities do not vary with δ . Therefore, in this case, theory gives us the point predictions: for $n = 3$, $k = 2$, and $\delta = 0.5$, theory predicts the vector of switching probabilities $(p, q_0, q_1) = (0.748, 1, 0.5)$. for $n = 3$, $k = 2$, and $\delta = 6$, theory predicts the vector of switching probabilities $(p, q_0, q_1) = (0, 1, 0.5)$.

Inspired by these point predictions, we formulate the following comparative statics hypotheses.

HYPOTHESIS 1: *Keeping the discount factor fixed, the second-round switching probability p is decreasing in the number of players n .*

This main hypothesis captures the idea that learning individuals in an organization account for the simultaneous learning of others by switching (exploring) less often as the team size grows.

HYPOTHESIS 2: *Fixing the number of players at either $n = 2$ or $n = 3$, the second-round switching probability p is decreasing in the discount factor δ .*

This second hypothesis tests whether variations in the payoff to achieving a success affect subjects' behavior; if the future matters little, then players should switch immediately (in round 2) but otherwise it may pay to wait (until round 3).

HYPOTHESIS 3: *The conditional switching probabilities in round 3 (q_0, q_1) should be invariant to changes in $n > 1$ or the discount factor.*

This third hypothesis is both a counterpoint to Hypothesis 2 and a test of individual rationality.

IV. Experimental Findings

In this section we report our main experimental findings and relate them to the theoretical predictions.⁸ We start by examining the aggregate data on second-round switching probabilities—the p 's. This addresses our central comparative statics hypothesis stated in the previous section that the second-round switching probability is decreasing in organizational size as well as in the discount factor. We then proceed with reporting the session-level data for the two-agent, three-agent, and individual-agent treatments. Here, in addition to the p 's, we report the conditional, third-round, switching probabilities, i.e., the probability of switching in round 3 conditional on not having switched in round 2, q_0 , and the probability of switching in round 3 conditional on having switched in round 2, q_1 . For each number $n = 1, 2, 3$ of agents, we test for differences in the p 's between treatments with different discount factors and for differences between q_0 and q_1 . Then, fixing the discount factor, we test for differences in the p 's and q 's between treatments. We then examine aggregate behavior over time and ask whether behavior changes dramatically with changes in the discount factor and whether there is an order effect depending on the sequence of change in the discount factor, high-low or low-high, in a session. Finally, we briefly examine individual behavior.

A. Aggregate Findings

Figure 1 shows the second-round switching probabilities for each treatment. The top panel displays the round 2 switching probabilities predicted by theory and the mean round 2 switching probabilities of the subjects for the case where $\delta = 0.5$ for individuals, pairs, and triples. The bottom panel does the same for $\delta = 6$. The data averages reported in Figure 1 are over all eight periods of all four sessions of a given treatment.

While the means of the observed round 2 switching probabilities do not coincide with the theoretical predictions, the comparative statics predictions of the theory *are* supported by the data. Our *first main experimental finding* is that, in line with our comparative statics hypotheses, the observed round 2 switching probability decreases as the size of the organization increases. When $\delta = 0.5$, the respective mean switching probabilities are 87.5 percent for singles (individuals), 63.5 percent for pairs, and 52.2 percent for triples. When $\delta = 6$, the mean switching probability falls from 79.9 percent in the individual case to 42.7 percent in the pairs case and 32.0 percent in the triples case.

The results presented in Figure 1 are disaggregated by experimental session in Table 2. Using the four session-level means for each treatment, a robust rank order test confirms that round 2 switching probabilities are significantly lower in the pairs treatment than in the singles treatment ($p \leq 0.05$) when $\delta = 0.5$ or when $\delta = 6$.⁹ Similarly, using the session-level data we also find that the round 2 switching probabilities are significantly lower in the triples treatment than in the pairs treatment ($p \leq 0.05$) when $\delta = 0.5$ or when $\delta = 6$. This evidence suggests that participants take into account the confounding effects of simultaneous learning.

Our *second main experimental finding* is that, in line with our comparative statics hypotheses in the pairs and triples treatments, the observed second-round switching probabilities decrease as we induce participants to be more patient by changing the discount factor from 0.5 to 6. For the pairs treatment, the predicted probability of switching in the second round is 80 percent when

⁸ All data from our experimental sessions are provided in the web Appendix.

⁹ For a description of the robust rank order test, see Sidney Siegel and N. John Castellan Jr. (1988). We used the robust rank order test for the other hypotheses tested in this section.

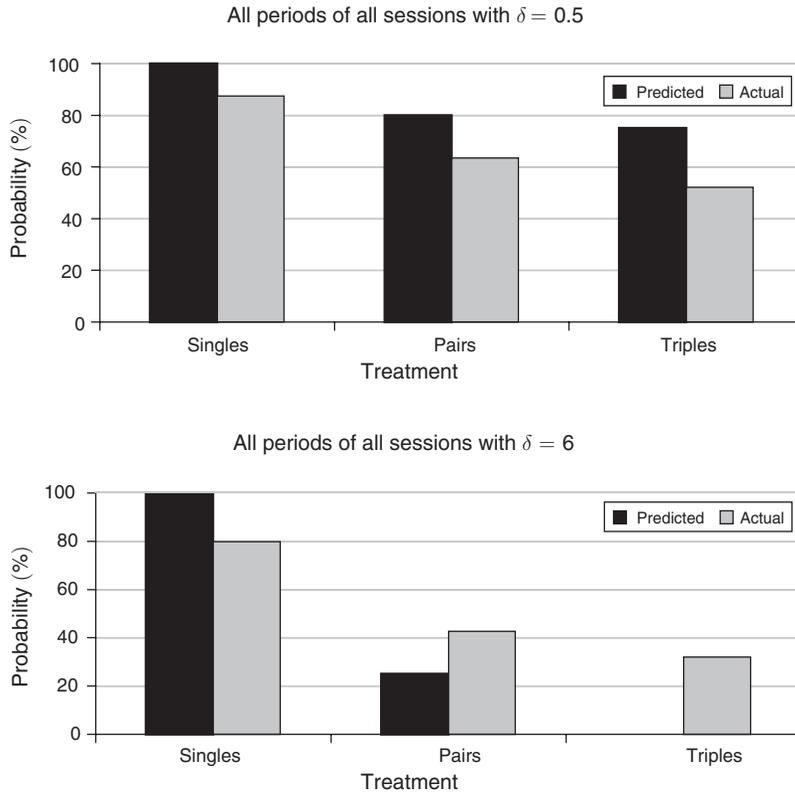


FIGURE 1. PREDICTED AND MEAN ROUND 2 SWITCHING PROBABILITIES OVER ALL PERIODS AND SESSIONS OF A GIVEN TREATMENT $\delta = 6$

$\delta = 0.5$ and falls to 25 percent when $\delta = 6$. In the experimental data, the mean second-round switching probability is 63.5 percent when $\delta = 6$ and falls to 42.7 percent when $\delta = 0.5$. Using the four session-level means for each treatment given in Table 1, this difference is significant ($p \leq 0.05$). Similarly, for the triples treatment, the predicted probability of switching in the second round is 75 percent when $\delta = 0.5$ and falls to 0 percent when $\delta = 6$. In the experimental data, the mean second-round switching probability is 52.2 percent when $\delta = 6$ and falls to 32 percent when $\delta = 0.5$. This difference is also significant using the session-level data. ($p \leq 0.05$).

The *third main experimental finding* concerns the conditional probabilities of switching in round 3 in the pairs and triples treatments. Recall that, regardless of the value of the discount factor or whether players are matched in pairs or triples, the probability of switching in round 3 conditional on not having switched in round 2 is $q_0 = 1.0$, while the probability of switching in round 3 conditional on having switched in round 2, $q_1 = 0.5$. The mean observed values of q_0 and q_1 in the four sessions conducted of the pairs and triples treatments are reported in Table 3.

Again, we see that while the point predictions of the theory are not borne out in the experimental data, the comparative static predictions of the theory find strong support. In particular, the conditional probability q_0 is greater than q_1 in both the pairs and the triples treatments. Using the four session-level means for q_0 and the corresponding session-level means for q_1 for the same treatment conditions (same δ , n), we can always reject the null hypothesis that $q_0 = q_1$ in favor

TABLE 2—PREDICTED AND OBSERVED MEAN ROUND-2 SWITCHING PROBABILITIES ACROSS TREATMENTS

δ	Grouping (n)	Predicted	Session 1	Session 2	Session 3	Session 4	Mean all
0.5	Singles	1.000	0.940	0.955	0.781	0.825	0.875
0.5	Pairs	0.800	0.657	0.561	0.555	0.767	0.635
0.5	Triples	0.750	0.465	0.486	0.514	0.622	0.522
6.0	Singles	1.000	0.816	0.861	0.761	0.759	0.799
6.0	Pairs	0.250	0.454	0.417	0.483	0.352	0.427
6.0	Triples	0.000	0.302	0.438	0.221	0.316	0.320

TABLE 3—PREDICTED AND OBSERVED MEAN CONDITIONAL ROUND-3 SWITCHING PROBABILITIES ACROSS PAIRS AND TRIPLES TREATMENTS

Pr.	δ	Grouping (n)	Predicted	Session 1	Session 2	Session 3	Session 4	Mean all
q_0	0.5	Pairs	1.0	0.568	0.556	0.518	0.497	0.535
q_0	0.5	Triples	1.0	0.505	0.397	0.558	0.572	0.508
q_0	6.0	Pairs	1.0	0.700	0.549	0.628	0.672	0.637
q_0	6.0	Triples	1.0	0.583	0.577	0.686	0.666	0.628
q_1	0.5	Pairs	0.5	0.272	0.401	0.313	0.314	0.325
q_1	0.5	Triples	0.5	0.299	0.357	0.401	0.300	0.339
q_1	6.0	Pairs	0.5	0.447	0.260	0.346	0.372	0.356
q_1	6.0	Triples	0.5	0.298	0.303	0.290	0.332	0.306

of the alternative that $q_0 > q_1$ ($p \leq 0.05$). This observed difference in switching probabilities conditional on having switched versus not having switched before (in both the pairs and triples treatments) is our third main experimental finding. It supports the theoretical prediction that the third-round switching probability when both actions have been taken previously is lower than when only one action has been taken previously.

Our *fourth main experimental finding* concerns the effect of varying organization size (n) on third-round switching probabilities. The third-round switching probabilities do not appear to change as the size of the organization increases from two to three members. Specifically, if we compare the four session mean values for q_0 in the case of pairs (as reported in Table 3) with the corresponding session mean values for q_0 in the case of triples (for the same value of $\delta = 0.5$ or 6) we are unable to reject the null hypotheses that the means come from the same distribution (both tests yield p -values > 0.10). Similarly, if we compare the four session-level mean values for q_1 in the case of pairs with the corresponding session mean values for q_1 in the case of triples (for the same value of $\delta = 0.5$ or 6) we are also unable to reject the null hypotheses that the means come from the same distribution ($p > 0.10$ for both tests). This finding supports the theoretical prediction that round 3-behavior should be invariant with respect to the size of the organization, i.e., whether $n = 2$ or $n = 3$.

Our *fifth main experimental finding* concerns the effect of varying the discount factor on third-round switching probabilities. Support for the theoretical prediction that round 3 conditional probabilities are invariant to changes in the discount factor is mixed. On the one hand, comparing the four session mean values of q_0 when $\delta = 0.5$ with the corresponding values of q_0 when $\delta = 6$ (holding n fixed at either 2 or 3), we can reject the null hypothesis that the q_0 values come from the same distribution in favor of the alternative that q_0 is higher when $\delta = 6$ than when $\delta = 0.5$ ($p \leq 0.05$ for both tests). On the other hand, comparing the four session mean values of q_1 when $\delta = 0.5$ with the corresponding values of q_1 when $\delta = 6$ (holding n fixed at either 2 or 3), we cannot reject the null hypothesis that these q_1 values come from the same distribution

($p > 0.10$). The higher value of q_0 when $\delta = 6$ than when $\delta = 0.5$ is likely due to the difference in monetary rewards that players could earn by achieving the success outcome in round 3 (\$6 versus \$0.50); this difference may have made it more salient to players who had not switched in round 2 (and who had not yet achieved a success) that their best response was to switch in round 3.

Our *sixth main experimental finding* is that in the singles treatments, somewhat anomalously, individuals sometimes fail to switch even though it is optimal to switch with probability 1 in *both* rounds 2 and 3 to one of the remaining unexplored cells. The mean round 2 switching probabilities from the 4 individual-treatment sessions were given earlier in Table 2. There we see that when $\delta = 0.5$, the average round 2 switching probability was 0.875 and when $\delta = 6$ it is a little lower, 0.799. While these data depart from the theoretical point prediction of 1.0, the null hypothesis of no difference between the two treatments in the round 2 switching probabilities cannot be rejected using the four session-level observations ($p > 0.05$). This behavioral anomaly of not switching all the time in the singles treatments may help explain the departures from the point predictions in the pairs and triples treatments. It suggests that for some individuals the choice of switching probability is not guided by rational deliberation. If these individuals randomly decide whether or not to switch, with no bias either way, the observed second-round switching probabilities in the pairs and triples treatments would be biased away from the point predictions in the direction of switching with probability 0.5, and this is what we see in the data.

The round 3 switching probabilities in the singles treatment are given in Table 4. For comparison purposes, Table 4 provides the same conditional round 3 switching probabilities that were reported and examined in the pairs and triples treatment (cf. Table 3) even though in the singles treatment, the prediction is that $q_0 = q_1 = 1.0$, i.e., the round 3 switching probability is *not* conditional on whether a switch was made in round 2. We therefore also report in Table 4 the *unconditional* round 3 switching probability, q .¹⁰ Whether we look at the mean unconditional or the mean conditional probabilities, there is no significant difference in the session-level means for q , q_0 , or q_1 as δ is increased from 0.5 to 6 ($p > 0.10$), consistent with theoretical predictions. On the other hand, holding δ fixed at 0.5 or 6, we can reject the null of no difference in the *conditional* probabilities in favor of the alternative that $q_1 > q_0$ ($p \leq 0.05$). Interestingly, while in our data for the pairs and triples treatments, second- and third-round switching probabilities are negatively correlated as theory predicts, in the singles treatment, these probabilities are positively correlated. It appears that, in the singles treatment, having made one irrational decision increases the probability of making another! However, caution is warranted in making much of this finding, as the percentage of the ten players who chose not to switch in rounds 2 and/or 3 of the singles treatment is always rather small (less than 20 percent of subjects on average).

Summarizing, our analysis of the session-level mean switching probabilities provides clear support for the claim that learning in organizations is quite different from individual learning; it seems that in organizations players *do* take account of the fact that other players are learning and adjust their probability of switching actions relative to the individual decision-making environment. They also appear responsive to changes in the discount factor. These results suggest that the assumption of fully rational, strategic agents provides a reasonable benchmark for modeling decentralized organizational learning in contrast to the approach that has been traditionally taken, as noted in the introduction.

¹⁰ We have verified that subjects in the singles treatment were switching in round 3 to profiles they had *not* previously played, e.g., if an individual played XX in round 1, and switched to XY in round 2, then a "switch" in round 3 was either to YY or YX , and not back to XX . This was the case in 99.8 percent of all reported instances of switching in round 3 of the singles treatment. This issue arises only in the singles treatment, where subjects have four rather than two choices.

TABLE 4—PREDICTED AND OBSERVED MEAN CONDITIONAL AND UNCONDITIONAL ROUND-3 SWITCHING PROBABILITIES IN THE SINGLES TREATMENTS

Pr.	δ	Predicted	Session 1	Session 2	Session 3	Session 4	Mean all
q_0	0.5	1.0	0.936	0.750	0.692	0.354	0.683
q_1	0.5	1.0	0.947	0.955	0.880	0.937	0.930
q (uncond.)	0.5	1.0	0.943	0.944	0.841	0.780	0.877
q_0	6.0	1.0	1.000	0.6890	0.739	0.442	0.717
q_1	6.0	1.0	1.000	0.9400	0.849	0.975	0.941
q (uncond.)	6.0	1.0	1.000	0.889	0.804	0.788	0.870

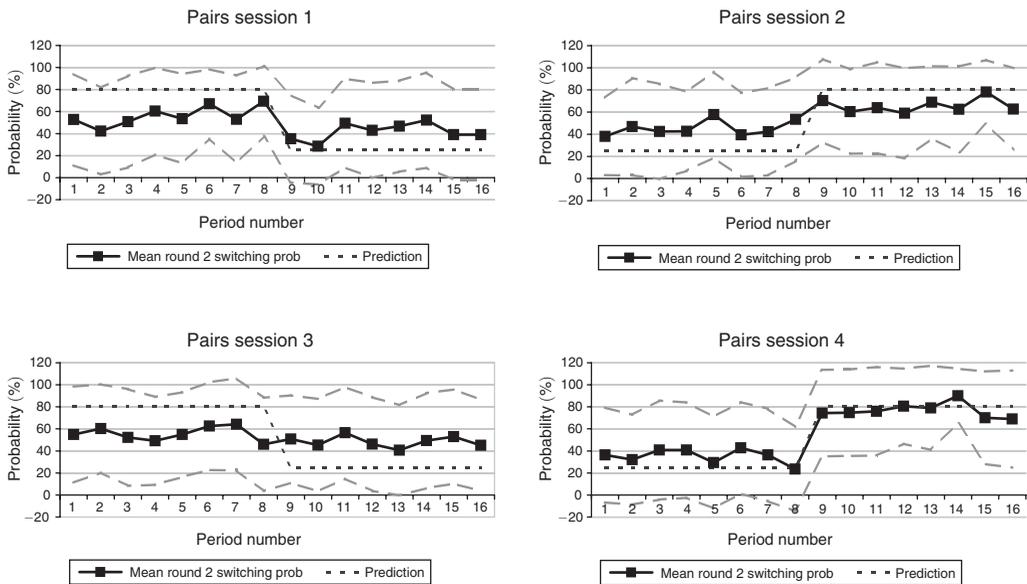


FIGURE 2. PREDICTED AND MEAN ROUND 2 SWITCHING PROBABILITIES WITH ONE STANDARD DEVIATION BOUND IN FOUR PAIRS SESSIONS

B. Behavior over Time

Thus far we have reported on session-level mean observations across treatments. In this section we provide a brief characterization of aggregate behavior over time.

Our first finding is that, in the pairs and triples treatments, we often observe a sharp jump in the aggregate round 2 switching probabilities when δ is switched from 0.5 to 6 (low-high treatment) or from 6 to 0.5 (high-low treatment), which is consistent with theoretical predictions. Figures 2 and 3 show the evolution of the mean round 2 switching probability over periods 1–16 for the four pairs and four triples sessions, respectively. In these figures, sessions 1 and 3 involve the high-low treatment and sessions 2 and 4 involve the low-high treatment. In Figure 2 we observe that in three out of four sessions, the mean value of p changes in the predicted direction at the time that δ is changed. The one exception is pairs session number 3, where the predicted decline in round 2 switching probabilities beginning with period 9 is not immediately apparent. In Figure 3, we observe that in all four sessions, the mean value of p changes in the predicted direction at the time δ is changed.

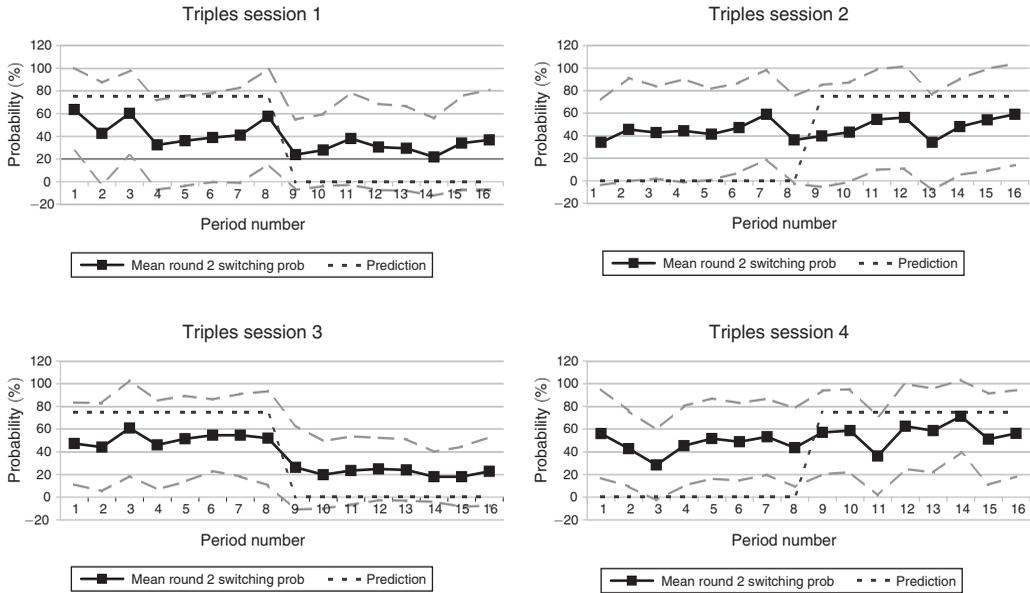


FIGURE 3. PREDICTED AND MEAN ROUND 2 SWITCHING PROBABILITIES WITH ONE STANDARD DEVIATION BOUND IN FOUR TRIPLES SESSIONS

A second finding is that there appears to be little evidence of any significant order effects in the pairs treatments. In particular, the sequence of discount factors we use, low-high (as in sessions 1 and 3) or high-low (as in sessions 2 and 4), does not appear to bias our findings away from the equilibrium predictions in any systematic fashion that is apparent in Figure 2. By contrast, in the triples treatment, there appears to be some evidence that the sequence of discount factor choices may matter for how closely aggregate behavior adheres to the comparative static predictions of the theory. In particular, in the low-high triples sessions (numbers 1 and 3), consistent with the theory, there is a marked drop-off in the round 2 switching probabilities from the first to the last eight periods of each session. In the high-low triples treatment, counter to the theory, there is much smaller change in the round 2 switching probabilities from the first to the last eight periods of each session.¹¹

Figure 4 shows the evolution of the mean round 2 switching probability over periods 1–16 in the four singles sessions. In this treatment, there should be no change in round 2 switching behavior when δ is changed, as the optimal strategy calls on players to switch with probability 1.0 regardless of the value of δ . Here again, sessions 1 and 3 involved the low-high treatment while sessions 2 and 4 involved the high-low treatment. In two of the four sessions, numbers 1 and 2, there is a small change in the mean value of p following the change in the discount factor, but it does not appear to be sustained beyond a couple of periods. Such changes are less clear in the other two sessions, numbers 3 and 4. This is consistent with the discussion surrounding Table 2, where we found that the session-level mean values for p in the singles treatment when $\delta = 0.5$ were not different from the session-level mean values for p when $\delta = 6$.

¹¹ A more formal test of whether the sequence of discount factors matters in the triples treatment would require more than the two observations that we have on the low-high and high-low treatments.

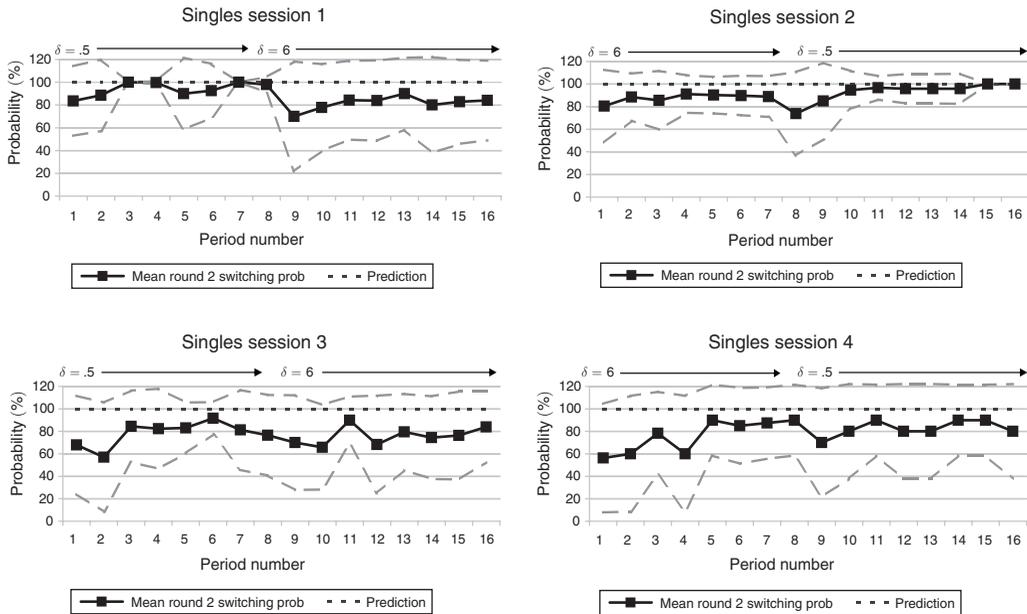


FIGURE 4. PREDICTED AND MEAN ROUND 2 SWITCHING PROBABILITIES WITH ONE STANDARD DEVIATION BOUND IN FOUR SINGLES SESSIONS

With regard to the behavior of the mean values of q_0 and q_1 , the prediction of the theory is that these probabilities should remain invariant over time to changes in the discount factor in all treatments. In the pairs and triples treatments, we should see $q_0 = 1$ and $q_1 = 0.5$, while in the singles treatment, $q_0 = q_1 = q = 1.0$. Consistent with the discussion of Table 3, we do not observe significant or sustained changes in the values of these conditional probabilities when the discount factor changes; that is, the graphs look similar to the those for p in the singles case.

Summarizing, the time path of the mean switching probabilities is sometimes volatile, but appears to be broadly consistent with the predictions of the theory. In particular, there appears to be a marked change in the round 2 switching probabilities immediately following a change in the discount factor in most of the pairs and triples sessions. By contrast, there appears to be little change in the round 2 switching probabilities in the singles treatments or in the round 3 conditional switching probabilities following a change in the discount factor, which is consistent with theoretical predictions.

V. Behavioral Models of Decision Making

Thus far, our analysis of the experimental subjects' switching behavior has been with reference to the Nash equilibrium point predictions under the maintained assumption of perfect rationality. In this section we relax the latter assumption and consider the predictions of three alternative models of boundedly rational decision making. In particular, we consider: an equilibrium-plus-noise model; a more sophisticated, stochastic equilibrium model of decision making known as quantal response equilibrium (QRE) (see, e.g., Richard D. McKelvey and Thomas R. Palfrey 1995, 1998); and, finally, a nonequilibrium model of decision making known as "step-" or "level-k" reasoning (see, e.g., Dale O. Stahl and Paul W. Wilson 1994, 1995; Rosemarie Nagel 1995; Teck-Hua Ho, Camerer, and Keith Weigelt 1998; Miguel A. Costa-Gomes, Crawford, and Bruno Broseta 2001; Costa-Gomes and Crawford 2006; and Crawford and Nagore Iriberry 2007).

A. Equilibrium Plus Noise

Perhaps the simplest way of reconciling the model with the fact that there is noise in the data is to consider a variant in which the predicted round 2 switching probability, $p(\eta)$, is a weighted average of the theoretical, equilibrium prediction p , and a purely random switching probability of $1/2$, that is, the *equilibrium-plus-noise model* specifies that subjects switch in round 2 with probability:

$$p(\eta) = \eta p + (1 - \eta) \frac{1}{2},$$

where $\eta \in [0,1]$. If we further impose the (sensible) restriction that the weight assigned to equilibrium, η , is the same for the conditional and unconditional switching probabilities of round 3, we can estimate η using the method of maximum likelihood. Under the equilibrium-plus-noise model, the predicted round 3 switching probability following no switch in round 2 is

$$q_0(\eta) = \eta q_0 + (1 - \eta) \frac{1}{2} = \eta + (1 - \eta) \frac{1}{2} = \frac{1 + \eta}{2},$$

and the predicted round 3 switching probability following a switch in round 2 is

$$q_1(\eta) = \eta q_1 + (1 - \eta) \frac{1}{2} = \eta \frac{1}{2} + (1 - \eta) \frac{1}{2} = \frac{1}{2}.$$

Since $q_1(\eta)$ does not depend on the parameter η , the corresponding factor in the likelihood function is constant and can be ignored when maximizing the likelihood function.

To construct the likelihood function, let ω be the total number of round 2 observations, ω_0 the total number of round 3 observations following “no switch and no success in round 2,” and ω_1 the total number of round 3 observations following a “switch in round 2, but no success in round 2.” Then, denote by σ the actual number of round 2 switches, σ_0 the actual number of round 3 switches following “no switch and no success” in round 2, and σ_1 the actual number of round 3 switches following a “switch in round 2, but no success in round 2.” The likelihood function is then proportional to

$$\tilde{L}(\eta) = p(\eta)^\sigma (1 - p(\eta))^{(\omega - \sigma)} q_0(\eta)^{\sigma_0} (1 - q_0(\eta))^{(\omega_0 - \sigma_0)} q_1(\eta)^{\sigma_1} (1 - q_1(\eta))^{(\omega_1 - \sigma_1)}.$$

Since we can ignore the term that involves $q_1(\eta)$, we instead maximize

$$L(\eta) = p(\eta)^\sigma (1 - p(\eta))^{(\omega - \sigma)} q_0(\eta)^{\sigma_0} (1 - q_0(\eta))^{(\omega_0 - \sigma_0)},$$

using pooled data from all sessions of a given treatment.

The results of this maximum likelihood (ML) estimation are given in Tables 5 and 6, which report the estimated values \hat{p} , \hat{q}_0 , and $\hat{\eta}$ using data from all eight periods, or for the first four or the last four periods of all sessions of a treatment (δ value) in the pairs (Table 5) and triples (Table 6) cases. Also reported are the results of a likelihood ratio test that compares the likelihood function for the unrestricted ML estimator for the equilibrium-plus-noise model with the likelihood function from a purely random version of that model, where we impose the restriction that $\eta = 0$ and use the same number of observations as for the unrestricted model. Specifically, the last two columns of Tables 5 and 6 report the likelihood ratio (LR) test statistic, (LR Stat $\equiv -2\ln\ell$, where ℓ is the ratio of the unrestricted to the restricted likelihood functions) and the associated p -value. The LR test statistic has a χ^2 distribution with degrees of freedom equal to the number of restrictions, in this case, one.

TABLE 5—EQUILIBRIUM-PLUS-NOISE MODEL, PAIRS TREATMENT: MAXIMUM LIKELIHOOD ESTIMATES AND LIKELIHOOD RATIO TEST RESULTS

Treatment: Pairs	\hat{p}	\hat{q}_0	$\hat{\eta}$	LR test statistic	p -value
All periods, $\delta = 6$	0.427	0.646	0.292	42.239	$p < 0.001$
First four periods, $\delta = 6$	0.429	0.644	0.287	20.575	$p < 0.001$
Last four periods, $\delta = 6$	0.426	0.649	0.298	21.680	$p < 0.001$
All periods, $\delta = 0.5$	0.584	0.640	0.280	34.252	$p < 0.001$
First four periods, $\delta = 0.5$	0.563	0.604	0.208	9.705	$p = 0.002$
Last four periods, $\delta = 0.5$	0.607	0.678	0.356	27.111	$p < 0.001$

TABLE 6—EQUILIBRIUM-PLUS-NOISE MODEL, TRIPLES TREATMENT: MAXIMUM LIKELIHOOD ESTIMATES AND LIKELIHOOD RATIO TEST RESULTS

Treatment: Triples	\hat{p}	\hat{q}_0	$\hat{\eta}$	LR test statistic	p -value
All periods, $\delta = 6$	0.359	0.642	0.283	61.395	$p < 0.001$
First four periods, $\delta = 6$	0.350	0.651	0.301	34.321	$p < 0.001$
Last four periods, $\delta = 6$	0.367	0.633	0.266	27.327	$p < 0.001$
All periods, $\delta = 0.5$	0.501	0.502	0.003	0.004	$p = 0.950$
First four periods, $\delta = 0.5$	0.515	0.530	0.059	0.519	$p = 0.471$
Last four periods, $\delta = 0.5$	0.487	0.474	-0.053	0.408	$p = 0.523$

Notice, first, that according to the LR test, the equilibrium-plus-noise model outperforms the purely random switching model in three of the four treatments. The one exception occurs for the triples $\delta = 0.5$ treatment, where we cannot reject the null hypothesis of no difference at any conventional level of significance. Notice further that the comparative static implications of the theory find support in the equilibrium-plus-noise estimates: the estimated values of p are, in both the pairs and triples treatment, greater when $\delta = 0.5$ than when $\delta = 6$, and in the pairs treatment, there is not much difference in the estimated value of q_0 as δ is varied; by contrast, for the triples $\delta = 0.5$ treatment, there is some difference in estimates of q_0 with changes in δ ; q_0 is estimated to be lower when $\delta = 0.5$ than when $\delta = 6$ (a finding that is nevertheless consistent with the experimental data—see Table 3). Regarding the $\hat{\eta}$ estimates, we note two things. First, in our estimation of the equilibrium-plus-noise model, η was not restricted to lie in the interval $[0, 1]$. Nevertheless, with one exception, the maximum likelihood estimates for η *always* lie within this interval, giving some support to the notion that switching probabilities might be appropriately modeled as a mixture of equilibrium and noise. Not surprisingly, the exception occurs for the last four periods of the triples $\delta = 0.5$ treatment, an instance where, as noted above, the equilibrium-plus-noise model does no better than the random switching model. Notice, second, that the weight on equilibrium, the estimated value of η , is low, approximately 0.3 for the pairs treatment under $\delta = 0.5$ and $\delta = 6$ and for the triples treatment under $\delta = 6$, and it is near zero in the triples $\delta = 0.5$ case. This finding mainly confirms that the switching probabilities p and q_0 deviate from equilibrium *in the direction of 0.5* (i.e., random switching), a fact that can also be ascertained from looking at the actual aggregate switching frequencies relative to the equilibrium predictions as reported earlier in Tables 2 and 3.

Summarizing, the equilibrium-plus-noise model provides a simple measure of the “closeness” of the data to equilibrium predictions. We have found that this model is a better fit to the data than purely random decision making for three of our four treatments (n, δ) . Further, under this

view of behavior, most of our equilibrium comparative statics predictions continue to hold. On the other hand, the estimated weight given to the equilibrium prediction in this model is low, and essentially zero in the triples $\delta = 0.5$ treatment. An obvious difficulty with the equilibrium-plus-noise model is that it assumes that otherwise rational players ignore the irrationality of others; this problem is addressed by the QRE model considered in the next section.

B. Quantal Response Equilibrium

In a QRE, players do not play best responses to their beliefs; instead they play “noisy” or “stochastic” best responses, rationally taking into account the noise in other players’ strategies as well; the latter is what distinguishes QRE from the equilibrium-plus-noise model. In QRE, the noise in players’ strategies is assumed to follow a specific distribution; in the case of the logit choice rule that we use, the distribution is log Weibull. The resulting QRE is called a logit equilibrium. In this section we use the experimental data to estimate a logit equilibrium model using the method of maximum likelihood. In addition to considering how the switching probabilities implied by the logit model compare with the theoretical predictions in the pairs and triples cases, we will also test whether the logit model’s predictions differ from purely random switching behavior.

We impose attainability constraints, i.e., we ignore the label of the first-round choice and focus on equilibria that are symmetric across players. The QRE is obtained using expected payoffs and a logit-choice rule. Specifically, for a particular game (e.g., pairs or triples) and for a given value of λ , the degree of rationality parameter, we begin by writing a system of logit-choice equations:

$$p = \frac{e^{\lambda u_s}}{e^{\lambda u_s} + e^{\lambda u_n}},$$

$$q_0 = \frac{e^{\lambda u_s^0}}{e^{\lambda u_s^0} + e^{\lambda u_n^0}},$$

$$q_1 = \frac{e^{\lambda u_s^1}}{e^{\lambda u_s^1} + e^{\lambda u_n^1}},$$

where the u_i^j represent expected payoffs from action i , switching (s), or not switching (n) that may also condition on the state j ($0 =$ no prior switch, $1 =$ prior switch). These expected payoff values are defined in the Appendix. We write the solution of this system of equations as $(p(\lambda), q_0(\lambda), q_1(\lambda))$ and maximize the likelihood function over different values of λ .

As in the equilibrium-plus-noise model, ω denotes the total number of round 2 observations, ω_0 the total number of round 3 observations following “no switch and no success in round 2,” and ω_1 the total number of round 3 observations following a “switch in round 2 but no success in round 2.” As before, σ is the actual number of round 2 switches, σ_0 the actual number of round 3 switches following “no switch and no success in round 2,” and σ_1 the actual number of round 3 switches following a “switch in round 2 but no success in round 2.”

The likelihood function is proportional to

$$L(\lambda) = p(\lambda)^\sigma (1 - p(\lambda))^{(\omega - \sigma)} q_0(\lambda)^{\sigma_0} (1 - q_0(\lambda))^{(\omega_0 - \sigma_0)} q_1(\lambda)^{\sigma_1} (1 - q_1(\lambda))^{(\omega_1 - \sigma_1)}.$$

In maximizing this likelihood function, we again used pooled data from all sessions of a given treatment (n, δ) . The value of q_1 is 0.5 in any logit equilibrium of our model. We therefore

TABLE 7—QUANTAL RESPONSE EQUILIBRIUM MODEL, PAIRS TREATMENT: MAXIMUM LIKELIHOOD ESTIMATES AND LIKELIHOOD RATIO TEST RESULTS

Treatment: Pairs	\hat{p}	\hat{q}_0	$\hat{\lambda}$	LR test statistic	p -value
All periods, $\delta = 6$	0.476	0.667	0.449	35.352	$p < 0.001$
First four periods, $\delta = 6$	0.477	0.660	0.427	16.307	$p < 0.001$
Last four periods, $\delta = 6$	0.474	0.673	0.471	19.109	$p < 0.001$
All periods, $\delta = 0.5$	0.588	0.623	3.260	37.785	$p < 0.001$
First four periods, $\delta = 0.5$	0.570	0.592	2.414	11.602	$p < 0.001$
Last four periods, $\delta = 0.5$	0.606	0.659	4.258	28.301	$p < 0.001$

TABLE 8—QUANTAL RESPONSE EQUILIBRIUM MODEL, TRIPLES TREATMENT: MAXIMUM LIKELIHOOD ESTIMATES AND LIKELIHOOD RATIO TEST RESULTS

Treatment: Triples	\hat{p}	\hat{q}_0	$\hat{\lambda}$	LR test statistic	p -value
All periods, $\delta = 6$	0.449	0.448	-0.242	4.551	$p = 0.033$
First four periods, $\delta = 6$	0.436	0.437	-0.288	3.320	$p = 0.068$
Last four periods, $\delta = 6$	0.460	0.458	-0.197	1.465	$p = 0.226$
All periods, $\delta = 0.5$	0.505	0.503	0.178	0.062	$p = 0.803$
First four periods, $\delta = 0.5$	0.512	0.507	0.411	0.145	$p = 0.703$
Last four periods, $\delta = 0.5$	0.496	0.500	-0.015	0.002	$p = 0.964$

maximize a truncated likelihood function (just as we did in the case of the equilibrium-plus-noise model), which yields maximum likelihood estimates for p and q_0 only.

The maximum likelihood estimates \hat{p} , \hat{q}_0 , and $\hat{\lambda}$ for the pairs and triples treatments are shown in Tables 7–8 using data from all eight periods, or from the first four or the last four periods of all sessions of a treatment (δ value). We again report the results of a likelihood ratio test that compares the likelihood function from the unrestricted ML estimator of the QRE model with the likelihood function from a purely random switching version of that model, where λ was restricted to be zero and where we have used the same number of observations as for the unrestricted model. Specifically, we report the LR test statistic (which is distributed χ^2 with one degree of freedom) and the associated p -value.

Consider, first, the estimates for the pairs $\delta = 6$ or $\delta = 0.5$ treatments, as reported in Table 7. For these treatments, the LR test indicates we can reject H_0 (the null no difference between the QRE model and the restricted, random switching model). Consistent with the comparative static predictions of the theory and the data, the QRE estimates for p are lower in the pairs $\delta = 6$ treatment than in the pairs $\delta = 0.5$ treatment. Similarly, consistent with the theory and the data, the QRE estimates of q_0 in the pairs treatment do not change much as the value of δ changes. A comparison of the QRE pairs estimates with the actual mean frequencies (over all periods) for p and q_0 , as reported in Tables 2 and 3, suggests that subjects were close to playing according to the estimated QRE in the pairs treatment. We note further that the QRE estimates of p and q_0 for the pairs treatment are quite similar to the equilibrium-plus-noise model estimates for the pairs treatment (compare Tables 5 and 7). Finally, we note that while the estimates of λ in the pairs treatment display some variation, they are all greater than zero; larger, positive values for λ are associated with greater rationality in decision making, while a λ equal to zero indicates random decision making.

Consider next the QRE estimates for the triples $\delta = 6$ and $\delta = 0.5$ treatments as reported in Table 8. Consistent with the theory and the data, the QRE estimates for p are lower when

$\delta = 6$ than when $\delta = 0.5$. However, inconsistent with the theory and the data, the estimates of q_0 are lower when $\delta = 6$ than when $\delta = 0.5$. Furthermore, the LR test results suggest that for all samples of the triples $\delta = 0.5$ treatment, as well as for the last four periods of the triples $\delta = 6$ treatment, we cannot reject H_0 , that the restricted, random switching model provides as likely an explanation of the data as the QRE model. Consistent with the latter finding, notice that the estimates for λ in the triples treatment are much closer to zero than in the comparable pairs treatment and are often slightly negative; we did not restrict our estimates of the QRE model parameter λ , just as we did not restrict our estimates of η in the equilibrium-plus-noise model. (Recall a similar finding for the equilibrium-plus-noise model, where estimates of η were closer to zero or even negative in the triples $\delta = 0.5$ treatment). One interpretation of these low λ estimates is that subjects in the triples treatment found their decision problem more challenging than did subjects in the pairs treatment. Comparing the QRE estimates for the triples treatment with the equilibrium-plus-noise estimates for this same treatment, we observe that for the $\delta = 6$ case, the equilibrium-plus-noise estimates of p and q_0 are somewhat closer to the mean switching frequencies found in the data than are the QRE estimates, while for the triples $\delta = 0.5$ treatment, the opposite finding obtains. We conclude that QRE does not appear to offer any improvement over the equilibrium-plus-noise model in explaining the data from our experiment.

Summarizing, QRE is a generalization of Nash equilibrium to the case of noisy best responses where noise enters via a logistic-choice function transformation of theoretical expected payoffs, and players take account of this noise in formulating best responses. We found that the QRE estimates for our *pairs* treatment improve upon a model of purely random switching behavior. Also for the pairs treatment, the QRE estimates of p and q_0 are consistent with the comparative static implications of the theory and bear some resemblance to the actual switching frequencies observed in the data. However, the same cannot be said for QRE estimates using data from the triples treatment, where QRE estimates of q_0 are inconsistent with both the comparative static implications of the theory and the data.

C. Level-k Analysis

Finally, we consider the fit of a nonequilibrium behavioral model known as “level-k analysis” to our experimental data. We suppose there are three player types. The lowest, Level 0 (L_0) players, switch randomly in every round. The next highest, Level 1 (L_1) players, play a best response to the L_0 types, and the highest, Level 2 (L_2) players, play a best response to the L_1 types. Our restriction to just three level types is in line with prior findings. For instance, Camerer, Ho, and Juin-Kuan Chong (2004) find that an average level of 1.5 fits the data from many games using their closely related “cognitive hierarchy” model (in which L_2 player types best respond to a Poisson mix of L_1 and L_0 player types). Crawford and Iriberry (2007) restrict themselves to L_0 , L_1 , and L_2 players and find a higher proportion of L_1 than either L_2 or L_0 players. Costa-Gomes and Crawford (2006) find roughly twice as many L_1 types as L_2 types.

The predictions of level-k analysis in our search-for-success game are found by considering the payoffs earned by the various player types from following reduced normal form strategies. Table 9 shows the payoff from play of the reduced normal form set of pure strategies in the two-player game, Switch (S) or No Switch (N) in periods 2 and 3 against one another. For instance, the strategy pair SS (switch in round 2 and switch back in round 3) against itself has an expected payoff of $1/3$. Given that L_0 types randomize uniformly, it follows that for $\delta = 0.5$, L_1 players are indifferent between SS and SN , and strictly prefer these two strategies to all others. By contrast, L_2 players who best respond to L_1 players’ decision to switch in round 2 strictly prefer NS to all other strategies. For $\delta = 6$, if we once again assume that L_0 players randomize uniformly, then

TABLE 9—PAYOFFS FROM PLAY OF PURE STRATEGIES AGAINST ONE ANOTHER IN THE TWO-PLAYER GAME

		Strategies of player 2			
		SS	SN	NS	NN
Strategies of player 1	SS	$\frac{1}{3}$	$\frac{1}{3} + \delta \frac{1}{3}$	$\frac{1}{3} + \delta \frac{1}{3}$	$\frac{1}{3}$
	SN	$\frac{1}{3} + \delta \frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3} + \delta \frac{1}{3}$	$\frac{1}{3}$
	NS	$\frac{1}{3} + \delta \frac{1}{3}$	$\frac{1}{3} + \delta \frac{1}{3}$	$\delta \frac{1}{3}$	$\delta \frac{1}{3}$
	NN	$\frac{1}{3}$	$\frac{1}{3}$	$\delta \frac{1}{3}$	0

Note: Player strategies are represented by pairs of values, S = Switch, N = No Switch; e.g., “SN” is the strategy Switch in round 2, No Switch in round 3.

L_1 players strictly prefer *NS* to all other strategies, and given this behavior, L_2 players' are indifferent between either *SS* and *SN*, but strictly prefer these strategies to all others.

Notice that, while the level- k analysis yields ambiguous predictions concerning round 3 switching behavior, it yields an unambiguous prediction regarding round 2 behavior in the two-player game: L_1 types should switch in round 2 when $\delta = 0.5$ and should not switch in round 2 when $\delta = 6$, and L_2 types should follow the precise opposite strategy, not switching in round 2 when $\delta = 0.5$ and switching in round 2 when $\delta = 6$.

We can also apply level- k analysis to the three-player, two-success game (our triples treatment). The payoff matrix from all different combinations of reduced normal form strategies in our three-player search-for-success game with two successes is given in Table 10, which can be read the same way as Table 9, except that the columns now report pure strategy pairs for two of the three players. Analysis of these payoffs for the three-person game under level- k reasoning reveals that L_1 and L_2 types should behave in precisely the same manner as predicted for the two-player game.

Table 11 reports results from a level- k analysis of round 2 switching decisions in both the pairs and triples sessions; as the level- k predictions for round 3 are ambiguous, we use only round 2 data in our analysis. Assignment of level type 1 or type 2, L_1 or L_2 , was made as follows. As noted above, when $\delta = 0.5$ in both the pairs and triples treatments, L_1 players strictly prefer switching in round 2 to not switching, while L_2 players strictly prefer not switching in round 2 to switching. When $\delta = 6$ in both the pairs and triples treatments, these predictions are precisely reversed: L_1 players strictly prefer not switching in round 2 to switching, while L_2 players strictly prefer switching in round 2 to not switching. Thus, we labeled each subject a ‘1’ if their round 2 switching behavior was in accord with L_1 play and a ‘2’ if their round 2 behavior was in accord with L_2 play. Using the 16 numerical assignments for each subject (across both the eight periods of the $\delta = 0.5$ treatment and the eight periods of the $\delta = 6$ treatment), we used the *modal* assignment to identify each subject's type (either 1 or 2). In the event where the mode was indeterminate—exactly eight of an individual subject's choices were in accord with level 1 play while the other eight were in accord with level 2 play—we classified that subject as a level 0, L_0 type.

Using these assignment rules, the distribution of player types is as given in the first three columns of Table 11. A striking finding across all sessions of both treatments is that the ratio of L_1 to L_2 types is, on average, approximately 2 to 1, which, as noted earlier, is consistent with prior findings, e.g., Costa-Gomes and Crawford (2006). Further, we find that a nonnegligible fraction (an average of 17 percent) of our subjects can be typed as level L_0 .

TABLE 10—PAYOFFS FROM PLAY OF PURE STRATEGIES AGAINST ONE ANOTHER IN THE THREE-PLAYER GAME

		Strategies of players 2–3							
		SS SS	SS SN	SS NS	SS NN	SN SS	SN SN	SN NS	SN NN
Strategies of player 1	SS	$\frac{2}{7}$	*	*	$\frac{2}{7}$	*	*	*	*
	SN	*	*	*	*	*	$\frac{2}{7}$	*	$\frac{2}{7}$
	NS	*	*	*	*	*	*	*	*
	NN	$\frac{2}{7}$	*	*	$\frac{2}{7}$	*	$\frac{2}{7}$	*	$\frac{2}{7}$

		Strategies of players 2–3							
		NS SS	NS SN	NS NS	NS NN	NN SS	NN SN	NN NS	NN NN
Strategies of player 1	SS	*	*	*	*	$\frac{2}{7}$	*	*	$\frac{2}{7}$
	SN	*	*	*	*	*	$\frac{2}{7}$	*	$\frac{2}{7}$
	NS	*	*	$\delta\frac{2}{7}$	$\delta\frac{2}{7}$	*	*	$\delta\frac{2}{7}$	$\delta\frac{2}{7}$
	NN	*	*	$\delta\frac{2}{7}$	$\delta\frac{2}{7}$	$\frac{2}{7}$	$\frac{2}{7}$	$\delta\frac{2}{7}$	0

Notes: * = $(2/7) + (5/7)\delta(2/6)$. Player strategies are represented by pairs of values, S = Switch, N = No Switch, e.g., “NS” is the strategy No Switch in round 2, Switch in round 3.

TABLE 11—LEVEL-TYPE ANALYSIS OF DATA FROM ROUND 2 DECISIONS IN THE PAIRS AND TRIPLES TREATMENTS

Treatment-session	Percent subjects level 0	Percent subjects level 1	percent subjects level 2	Implied round 2 Sw freq $\delta = 0.5$	Actual round 2 Sw freq $\delta = 0.5$	Implied round 2 Sw freq $\delta = 6$	Actual round 2 Sw freq $\delta = 6$
Pairs-1	0.15	0.55	0.30	0.63	0.66	0.38	0.45
Pairs-2	0.15	0.70	0.15	0.78	0.56	0.23	0.42
Pairs-3	0.05	0.50	0.45	0.53	0.56	0.48	0.48
Pairs-4	0.20	0.75	0.05	0.85	0.77	0.15	0.35
Pairs 1-4	0.14	0.63	0.24	0.69	0.64	0.31	0.43
Triples-1	0.27	0.53	0.20	0.67	0.47	0.34	0.30
Triples-2	0.13	0.53	0.33	0.60	0.49	0.40	0.44
Triples-3	0.20	0.60	0.20	0.70	0.51	0.30	0.22
Triples-4	0.20	0.60	0.20	0.70	0.62	0.30	0.32
Triples 1-4	0.20	0.57	0.33	0.67	0.52	0.33	0.32

Using these frequencies for the three player types, we can obtain level-k predictions for the actual switching frequencies in the $\delta = 0.5$ and $\delta = 6$ versions of the pairs and triples treatments. Specifically, when $\delta = 0.5$, the frequency of switching in round 2 should equal the percentage of L_1 types plus 1/2 times the percentage of L_0 (random) types. This is computed in the fourth column of Table 11. When $\delta = 6$, the frequency of switching in round 2 should equal the percentage of L_2 types + 1/2 times the percentage of L_0 types; this prediction is given in the sixth column of Table 11. The fifth and seventh columns report the actual switching frequencies reproduced

from Table 2 for comparison purposes. The main finding here is that the correlation between the level- k predicted round 2 switching frequencies and the actual round 2 switching frequencies is strongly positive. In the pairs treatment, the correlation across both the $\delta = 0.5$ and $\delta = 6$ cases is .91, and in the triples treatment, the correlation across both the $\delta = 0.5$ and $\delta = 6$ cases is 0.70. The high correlation coefficient suggests that the level- k model has some predictive power for round 2 switching decisions in both the pairs and triples treatments.

Summarizing, the level- k predictions are unambiguous only for round 2 of the search-for-success game. Nevertheless, it appears that a level- k analysis provides a reasonable, nonequilibrium characterization of round 2 switching behavior. Level- k thinking is frequently declared particularly appropriate for explaining behavior in first encounters with a game. In our setting, it appears that it helps explain behavior equally well after repeated exposure to the game. This could be a result of the fact that players get very little feedback between instances of the search-for-success game. Another possible explanation is that in the search-for-success game, the level- k predictions do not differ much from the symmetric equilibrium prediction.

VI. Individual Behavior

Finally, it is of interest to consider the behavior of individual subjects. In this section we focus exclusively on individual switching behavior in round 2, as behavior in this round is predicted to vary with changes in the value of δ and the group size. We first ask whether the differences in round 2 switching frequencies across the two different δ values and group sizes (as reported in Table 2) reflect changes in each individual's behavior, or whether the aggregate differences are due only to a subset of individuals. Figure 5 provides striking evidence that the distribution of individual round 2 switching behavior is quite distinct across treatments. This figure shows the empirical (weighted) cumulative distribution of round 2 switching frequencies using data from all four sessions of the four treatments, i.e., for each switching frequency it reports the proportion of all individuals who switched with that frequency or with a lower frequency. Notice that the cumulative frequency distribution for the pairs $\delta = 0.5$ treatment first-order stochastically dominates that of the triples $\delta = 0.5$ treatment. The latter first-order stochastically dominates the pairs $\delta = 6$ treatment, which in turn first-order stochastically dominates the triples $\delta = 6$ treatment. This ordering is precisely in accord with the comparative static implication of the symmetric equilibrium predictions, which (you will recall) predict that round 2 switching frequencies should be 0.80, 0.75, 0.25, and 0 across these four treatments, respectively. A two-sample, two-sided Kolmogorov-Smirnov test confirms that round 2 switching frequencies are significantly greater in the pairs treatment when $\delta = 0.5$ ($\delta = 6$) relative to the triples treatment when $\delta = 0.5$ ($\delta = 6$) ($\text{Pr} \leq 0.025$ in both comparisons).

Second, we examine how individual behavior relates to the symmetric Nash equilibrium prediction. In addressing this question, we calculated each subject's round 2 switching frequency in the eight periods played under $\delta = 0.5$ and in the eight periods played under $\delta = 6$ in either the pairs or the triples treatments. Thus for each subject i , we have an average round 2 switching frequency pair $(p_i^{0.5}, p_i^6)$. In Figure 6 we use a bubble chart format to plot these frequency pairs using all data from the four sessions of a treatment (pairs, triples). The size of each bubble indicates the number of individual observations of that frequency pair; the smallest bubble corresponds to a single, individual observation. The top panel of Figure 6 shows the individual round 2 switching frequencies in the pairs treatment, while the bottom panel does the same for the triples treatment. We have also indicated in Figure 6 the location of the unique symmetric Nash equilibrium and we remind the reader of the overall average round 2 switching frequencies (averaging over all subjects) as previously reported in Table 2.

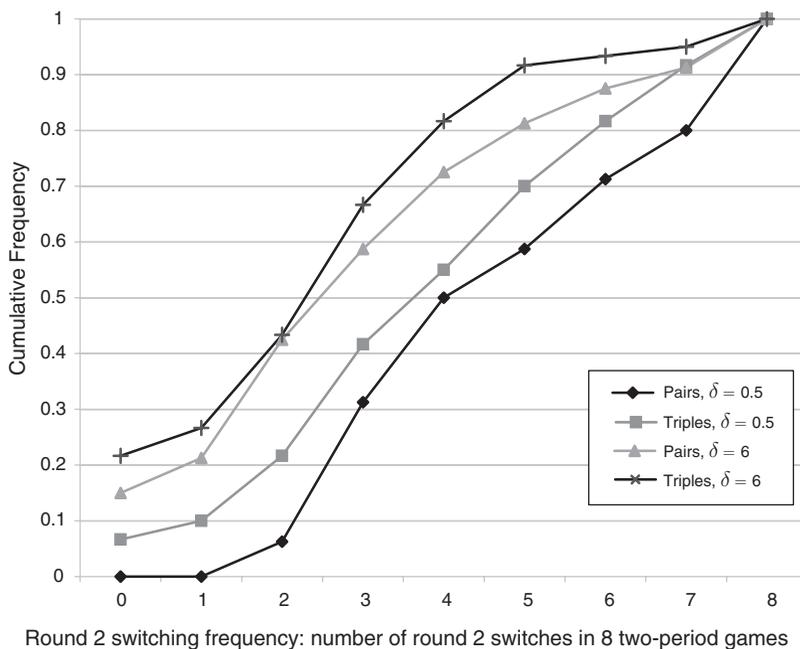


FIGURE 5. CUMULATIVE FREQUENCY DISTRIBUTION OF INDIVIDUAL ROUND 2 SWITCHING FREQUENCIES ACROSS TREATMENTS

Figure 6 reveals several interesting findings. First, there is considerable heterogeneity in the individual subjects' switching frequencies. Bear in mind, however, that in our environment, where individuals interact anonymously in a population setting, our equilibrium prediction *can* be realized as the population distribution over heterogeneous individual strategies, which themselves may be either pure or mixed (see, e.g., Ariel Rubinstein 1991).¹² Second, and more important, a majority of the mass of the individual frequencies lies below the 45 degree line in both panels of Figure 6, which is consistent with the location of the symmetric Nash equilibrium in both treatments (pairs, triples) and suggests that individual subjects were taking into account the impact on expected payoffs from a change in the value of δ . Finally, notice that in the pairs case, there is more mass associated with $p_i^{0.5} = 1$ than in the triples case and, symmetrically, in the triples case, there is more mass associated with $p_i^6 = 0$ than in the pairs case. Indeed, in the pairs treatment, a sizeable subset of subjects (16/80) appear to be following a heuristic switching strategy of the form $(p_i^{0.5}, p_i^6) = (1, x)$, where $x \in [0, 1]$. Similarly, a sizeable subset of subjects (13/60) in the triples treatment appears to have been using a heuristic switching strategy of the form $(p_i^{0.5}, p_i^6) = (x, 0)$, where $x \in [0, 1]$. These heuristics are roughly consistent with the play of level 1 (L_1) types; recall from our level-k analysis that in both the pairs and triples treatments, a L_1 type would switch in round 2 when $\delta = 0.5$ and not switch when $\delta = 6$, i.e., the L_1 strategy is $(p_i^{0.5}, p_i^6) = (1, 0)$. By contrast, there appear to be no *pure* level 2 (L_2) types, who would always play the strategy opposite to L_1 types, i.e., $(p_i^{0.5}, p_i^6) = (0, 1)$, though there is some mass above the 45 degree line in both the pairs and triples treatments. Finally, Figure 6 reveals evidence for level

¹² This parallels the Bayesian interpretation of Nash equilibrium that was articulated by John C. Harsanyi (1973) (using incomplete information) and later by Robert J. Aumann (1987) (without recourse to incomplete information), according to which a player's mixed strategy is an expression of the other players' ignorance of his (pure) strategy.

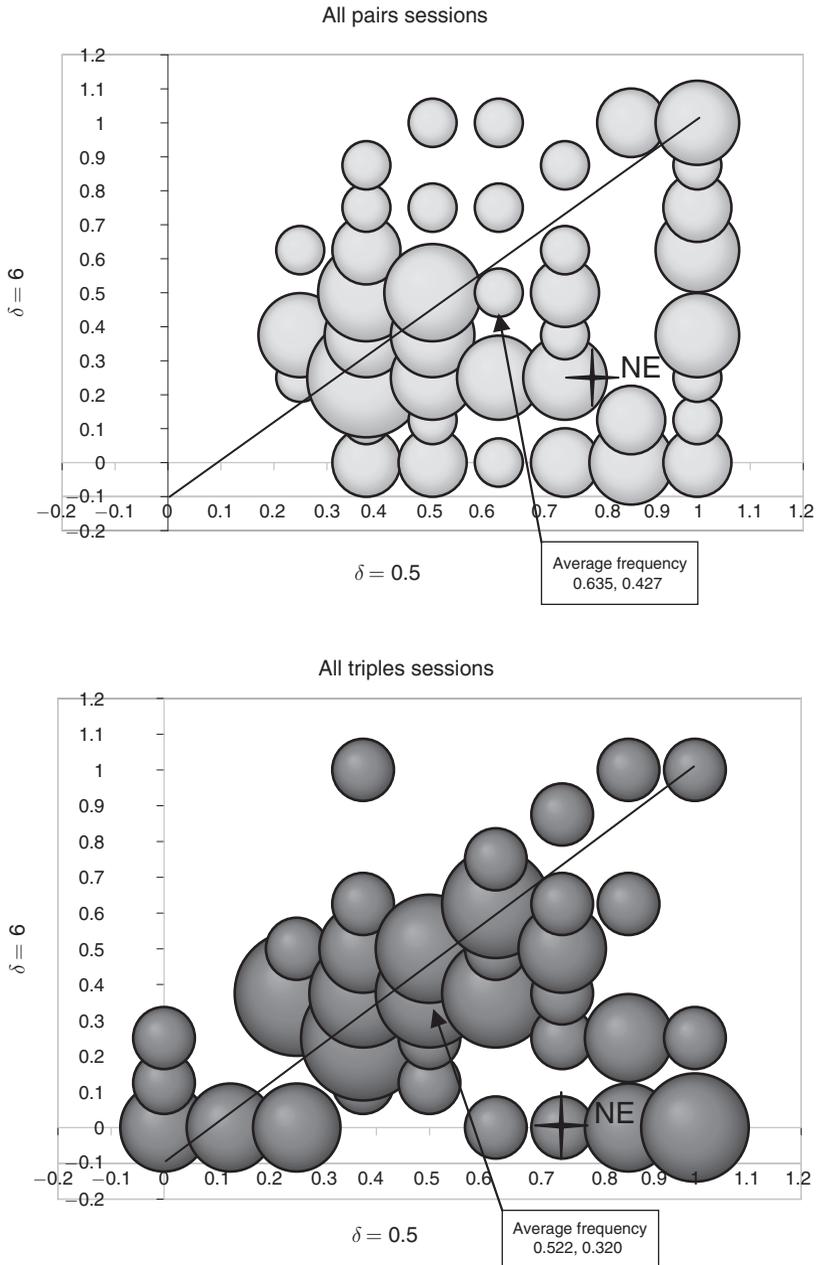


FIGURE 6. INDIVIDUAL SWITCHING FREQUENCIES AS A FUNCTION OF DELTA
 (Pooled data from all four Pairs (top) and Triples (bottom) sessions)

0 (L_0) types who switch with probability 0.5 for both values of δ , as evidenced by the mass at $(p_i^{\delta=0.5}, p_i^{\delta=0.6}) = (0.5, 0.5)$ in both the pairs and triples treatments.

VII. Conclusion

We have experimentally tested a stylized model of organizational learning that assumes rational agents, sparse information, and decentralized learning. In the random matching environment studied here, we find that the comparative static implications of the unique symmetric equilibrium are largely confirmed by our experimental data. The main qualitative insight from our experiment is that the rationality assumption has predictive power; the common practice in the organizational learning literature of assuming boundedly rational actors may be unwarranted. Specifically, we observe that team members cope with the confounding effects of the simultaneous learning of others by changing actions less frequently than individuals or teams of smaller size. Furthermore, we find that team members switch actions less frequently if the future becomes relatively more important. These are striking findings and to our knowledge they constitute the first direct evidence yet provided that learning individuals take into account the simultaneous learning of other agents.

We view these results as providing an important benchmark for future work. We plan to build on this investigation by changing the basic environment so that either tacit or explicit coordination of learning strategies becomes easier. As for tacit coordination, we expect that with fixed matchings and public information about the actions taken by team members, subjects would be more likely to develop *routines*, which would correspond to the asymmetric, but efficient, equilibria (where the number of strategy profiles explored in T periods was maximized). Explicit coordination via preplay communication similarly is likely to increase efficiency, and one could try to determine the relative value of “directives” (one-way communication) versus “committee meetings” (two-way communication).

APPENDIX: EXPECTED PAYOFF CALCULATIONS FOR QRE ESTIMATION

The quantal response equilibria of interest have the form (p, q_0, q_1) and are derived using expected payoffs and the logit choice rule. Here we provide details about the expected payoff calculations for the two- and three-player cases. The notation here is the same as in the paper.

For the two player case, define:

$$\begin{aligned} u_s &:= \text{payoff from switching in round 2 against } (p, q_0, q_1) \\ &= \frac{1}{3} + \frac{2}{3} \delta \left[p 2q_1(1 - q_1) \frac{1}{2} + (1 - p)q_0 \frac{1}{2} \right], \end{aligned}$$

$$\begin{aligned} u_n &:= \text{payoff from not switching in round 2 against } (p, q_0, q_1) \\ &= p \frac{1}{3} + \delta \left[p \frac{2}{3} q_0 \frac{1}{2} + (1 - p)(1 - (1 - q_0)^2) \frac{1}{3} \right], \end{aligned}$$

$$\begin{aligned} u_s^0 &:= \text{payoff from switching in round 3 after no switch in round 2 against } (p, q_0, q_1) \\ &= \delta \left[\tilde{p} \frac{1}{2} + (1 - \tilde{p}) \frac{1}{3} \right] \text{ where} \\ \tilde{p} &= \frac{\frac{2}{3}p}{\frac{2}{3}p + (1 - p)}, \end{aligned}$$

$$u_n^0 := \text{payoff from not switching in round 3 after no switch in round 2 against } (p, q_0, q_1)$$

$$= \delta \left[(1 - \tilde{p})q_0 \frac{1}{3} \right],$$

$$\begin{aligned} u_s^1 &:= \text{payoff from switching in round 3 after a switch in round 2 against } (p, q_0, q_1) \\ &= \delta \left[p(1 - q_1) \frac{1}{2} + (1 - p)q_0 \frac{1}{2} \right], \end{aligned}$$

$$\begin{aligned} u_n^1 &:= \text{payoff from not switching in round 3 after a switch in round 2 against } (p, q_0, q_1) \\ &= \delta \left[(1 - p)q_0 \frac{1}{2} + pq_1 \frac{1}{2} \right]. \end{aligned}$$

Similarly, for the three-player case, define:

$$\begin{aligned} u_s &:= \text{payoff from switching in round 2 against } (p, q_0, q_1) \\ &= \frac{2}{7} + \frac{5}{7} \delta \left\{ p^2 [3q_1(1 - q_1)^2 + 3q_1^2(1 - q_1)] \frac{2}{6} \right. \\ &\quad + 2p(1 - p) \left[q_0 \frac{2}{6} + (1 - q_0)(2q_1(1 - q_1)) \frac{2}{6} \right] \\ &\quad \left. + (1 - p)^2 [1 - (1 - q_0)^2] \frac{2}{6} \right\}, \end{aligned}$$

$$\begin{aligned} u_n &:= \text{payoff from not switching in round 2 against } (p, q_0, q_1) \\ &= 2p(1 - p) \left[\frac{2}{7} + \frac{5}{7} \delta (1 - (1 - q_0)^2) \frac{2}{6} \right] \\ &\quad + p^2 \left[\frac{2}{7} + \frac{5}{7} \delta \left(q_0 \frac{2}{6} + (1 - q_0)(2q_1(1 - q_1)) \frac{2}{6} \right) \right] \\ &\quad + (1 - p)^2 \left[\delta \frac{2}{7} (1 - (1 - q_0)^3) \right], \end{aligned}$$

$$\begin{aligned} u_s^0 &:= \text{payoff from switching in round 3 after no switch in round 2 against } (p, q_0, q_1) \\ &= \delta \left[\tilde{p} \frac{2}{6} + (1 - \tilde{p}) \frac{2}{7} \right] \text{ where} \\ &\quad \tilde{p} = \frac{(1 - (1 - p)^2)5/7}{(1 - (1 - p)^2)5/7 + (1 - p)^2}, \end{aligned}$$

$$\begin{aligned} u_n^0 &:= \text{payoff from not switching in round 3 after no switch in round 2 against } (p, q_0, q_1) \\ &= \delta \left\{ \frac{p^2 5/7}{p^2 5/7 + 2p(1 - p)5/7 + (1 - p)^2} (2q_1(1 - q_1)) \frac{2}{6} \right. \\ &\quad \left. + \frac{2p(1 - p) 5/7}{p^2 5/7 + 2p(1 - p)5/7 + (1 - p)^2} q_0 \frac{2}{6} \right\} \end{aligned}$$

$$+ \frac{(1-p)^2}{p^2 5/7 + 2p(1-p)5/7 + (1-p)^2} (1 - (1 - q_0)^2) \frac{2}{7} \Big\},$$

u_s^1 := payoff from switching in round 3 after a switch in round 2 against (p, q_0, q_1)

$$= \delta \left\{ p^2 (1 - q_1^2) \frac{2}{6} \right. \\ + 2p(1-p) \left(q_0 \frac{2}{6} + (1 - q_0)(1 - q_1) \frac{2}{6} \right) \\ \left. + (1-p)^2 (1 - (1 - q_0)^2) \frac{2}{6} \right\},$$

u_n^1 := payoff from not switching in round 3 after a switch in round 2 against (p, q_0, q_1)

$$= \delta \left\{ p^2 (1 - (1 - q_1)^2) \frac{2}{6} \right. \\ + 2p(1-p) \left(q_0 \frac{2}{6} + (1 - q_0) q_1 \frac{2}{6} \right) \\ \left. + (1-p)^2 (1 - (1 - q_0)^2) \frac{2}{6} \right\}.$$

REFERENCES

- Alpern, Steve, and Diane J. Reyniers.** 2002. "Spatial Dispersion as a Dynamic Coordination Problem." *Theory and Decision*, 53(1): 29–59.
- Argote, Linda.** 1999. *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Norwell, MA: Kluwer Academic Publishers.
- Aumann, Robert J.** 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica*, 55(1): 1–18.
- Banerjee, Abhijit V.** 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3): 797–817.
- Bhaskar, V.** 2000. "Egalitarianism and Efficiency in Repeated Symmetric Games." *Games and Economic Behavior*, 32(2): 247–62.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch.** 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades." *Journal of Political Economy*, 100(5): 992–1026.
- Blume, Andreas.** 2000. "Coordination and Learning with a Partial Language." *Journal of Economic Theory*, 95(1): 1–36.
- Blume, Andreas, and April M. Franco.** 2007. "Decentralized Learning from Failure." *Journal of Economic Theory*, 133(1): 504–23.
- Blume, Andreas, and Uri Gneezy.** 2000. "An Experimental Investigation of Optimal Learning in Coordination Games." *Journal of Economic Theory*, 90(1): 161–72.
- Blume, Andreas, and Uri Gneezy.** 2009. "Cognitive Forward Induction and Coordination without Common Knowledge: An Experimental Study." University of Pittsburgh Working Paper.
- Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle.** 1998. "Experimental Evidence on the Evolution of Meaning of Messages in Sender-Receiver Games." *American Economic Review*, 88(5): 1323–40.

- Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle.** 2001. "Evolution of Communication with Partial Common Interest." *Games and Economic Behavior*, 37(1): 79–120.
- Blume, Andreas, Douglas V. DeJong, George R. Neumann, and N. E. Savin.** 2002. "Learning and Communication in Sender-Receiver Games: An Econometric Investigation." *Journal of Applied Econometrics*, 17(3): 225–47.
- Bolton, Patrick, and Christopher Harris.** 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–74.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong.** 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics*, 119(3): 861–98.
- Christensen, Clayton M.** 1997. *The Innovator's Dilemma*. Boston: Harvard Business School Press.
- Costa-Gomes, Miguel A., and Vincent P. Crawford.** 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96(5): 1737–68.
- Costa-Gomes, Miguel A., Vincent P. Crawford, and Bruno Broseta.** 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 69(5): 1193–1235.
- Crawford, Vincent P., and Hans Haller.** 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games." *Econometrica*, 58(3): 571–95.
- Crawford, Vincent P., and Nagore Iriberri.** 2007. "Fatal Attraction: Saliency, Naivete, and Sophistication in Experimental 'Hide-and-Seek' Games." *American Economic Review*, 97(5): 1731–50.
- Cyert, Richard M., and James G. March.** 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice Hall.
- Day, Richard H., and E. Herbert Tinney.** 1968. "How to Co-operate in Business without Really Trying: A Learning Model of Decentralized Decision Making." *Journal of Political Economy*, 76(4): 583–600.
- Edmondson, Amy C.** 2004. "Learning from Failure in Health Care: Frequent Opportunities, Pervasive Barriers." *Quality and Safety in Health Care*, 13(6): 3–9.
- Harsanyi, John C.** 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of Game Theory*, 2(1): 1–23.
- Ho, Teck-Hua, Colin F. Camerer, and Keith Weigelt.** 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests.'" *American Economic Review*, 88(4): 947–69.
- Kramarz, Francis.** 1996. "Dynamic Focal Points in N-Person Coordination Games." *Theory and Decision*, 40(3): 277–313.
- Levinthal, Daniel A., and James G. March.** 1993. "The Myopia of Learning." *Strategic Management Journal*, 14(S2): 95–112.
- Lounamaa, Pertti H., and James G. March.** 1987. "Adaptive Coordination of a Learning Team." *Management Science*, 33(1): 107–123.
- Marschak, Jacob.** 1960. "Theory of an Efficient Several-Person Firm." *American Economic Review*, 50(2): 541–48.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10(1): 6–38.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1(1): 9–41.
- Nagel, Rosemarie.** 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85(5): 1313–26.
- North, Douglass C.** 1994. "Economic Performance through Time." *American Economic Review*, 84(3): 359–68.
- Rubinstein, Ariel.** 1991. "Comments on the Interpretation of Game Theory." *Econometrica*, 59(4): 909–24.
- Sen, Sandip, and Mahendra Sekaran.** 1998. "Individual Learning of Coordination Knowledge." *Journal of Experimental & Theoretical Artificial Intelligence*, 10(3): 333–56.
- Siegel, Sidney, and N. John Castellan, Jr.** 1988. *Nonparametric Statistics for the Behavioral Sciences*. 2nd ed. New York: McGraw Hill.
- Simon, Herbert A.** 1947. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. New York: Macmillan.
- Stahl, Dale O., and Paul W. Wilson.** 1994. "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior and Organization*, 25(3): 309–27.
- Stahl, Dale O., and Paul W. Wilson.** 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior*, 10(1): 218–54.
- Weber, Roberto A., and Colin F. Camerer.** 2003. "Cultural Conflict and Merger Failure: An Experimental Approach." *Management Science*, 49(4): 400–15.
- Wiseman, Thomas.** 2005. "A Partial Folk Theorem for Games with Unknown Payoff Distributions." *Econometrica*, 73(2): 629–45.