02/14/00 4:09 PM

Social Learning in a Social Hierarchy: An Experimental Study

by
Catherine C. Eckel
Department of Economics (0316)
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061
Email: eckelc@vt.edu

and
Rick K. Wilson
Department of Political Science
Rice University
Houston, TX

Abstract: Social learning occurs in the context of a structure of social interaction among agents. Models of the social structure typically do not include hierarchical elements, such as status differences among agents. In this paper we report the results of experiments designed to examine learning in the presence of a commonly-observed, higher-status agent. We test strategy choice by subjects in a 3x3 coordination game with two equilibria. We find that the commonly-observed agent can influence the choice between two equilibria, one of which has a higher payoff but is risky, and the other of which has a lower payoff but lower risk. However, there are limits to this agent's influence. A commonly-observed agent cannot move subjects away from an equilibrium that dominates the other on all dimensions, and he cannot induce subjects to choose a strategy that does not lead to an equilibrium. We explore the implications of this research for understanding the influence of high-status members of society.

*We thank Shabnam Modarres-Mousavi for research assistance, Cathleen A. Johnson for invaluable assistance in recruiting subjects and conducting the experiments at Virginia Tech, and Nancy Martorano for conducting experiments at Rice University. We also thank Megan Wilson for tracking down questions for the "trivia quiz" used in the experiment. The experiments at Virginia Tech were conducted at the Laboratory for the Study of Human Thought and Action. The paper benefited from comments of participants in the Workshop on Social Interaction, Santa Fe Institute, October, 1999. This research was supported by a grant from the National Science Foundation.

Social Learning in a Social Hierarchy: An Experimental Study

I. Introduction

Social learning occurs in the context of a structure of social interaction among agents. Models of the social structure typically do not include hierarchical elements, such as status differences among agents. Research in this area has neglected the possibility that social hierarchies may affect learning. The goal of our work is to examine learning in a laboratory experiment with a social hierarchy. In this paper we report the results of experiments designed to examine learning in the presence of a commonly-observed, higher-status agent.

Our interest in learning stems from the failure of game theory to predict the behavior of subjects in a variety of simple games in the laboratory. This failure has evoked two main responses from social scientists, both of which seek to incorporate elements of psychology into models of decision-making in games. The first approach argues that players' motivations are far richer than typical payoff-based utility functions will permit, and modifies the utility function of players to incorporate other-regarding preferences. The goal is to explain initial play, that is, how subjects formulate an initial strategy choice in the social environment presented in an experiment. Agents may care about the payoffs of others, their own relative earnings, or other aspects of the distribution of payoffs.

The second approach focuses on bounded rationality and suggests that subjects must first learn about an environment before they can be expected to choose strategies that serve their self-interest. This approach models how subjects learn, maintaining the traditional utility function. Many models employ evolutionary game theory, modeling agents as strategies, where the most successful strategies reproduce at a greater rate than others and less successful strategies die out. Several researchers have specified a structure for social interaction, which can be thought of as a social structure. In general, the social structure can have a strong impact on the equilibrium that the system or agents reaches.

We agree that both approaches are important. The decisions that agents make reflect utility functions with a social component, as well as a capacity for learning. We see learning itself as a social activity. Agents may care not only about the payoffs, actions, or characteristics of others in choosing a strategy, but also may care about the payoffs, actions or characteristics of others in choosing whom to observe and imitate. Models of social interaction and learning fail to take into account the possibility that agents may choose from whom to learn. Some agents in a social hierarchy may be focal, attracting greater attention than other agents. These focal agents may play a stronger role in determining the equilibrium that a system reaches.

In contrast to most theoretical research on learning, Bala and Goyal (1998) develop a network model with a commonly observed agent that they term the "royal family." The actions and payoffs of the royal family are observed by all other agents. They show that the choice of a product or technology with uncertain quality can be affected by the choices of the royal family, either positively or negatively.

Our experiment tests the effect of a high-status, commonly-observed agent on the choice of strategy in a coordination game. We adopt the 3x3 coordination games developed by Cooper, DeJong, Forsythe, and Ross (1990). Our results show that the royal family can affect the behavior of agents, but that the power of the commonly-observed agent is limited.

Section II discusses previous research on learning in experimental games; Section III surveys work (primarily from other fields) that leads us to believe that status will affect learning. Our experimental design and procedure are explained in sections IV and V, and VI contains results and analysis. A final section concludes the paper.

II. Modeling learning in games

Many researchers are using evolutionary game theory to model the dynamics of behavior in games. One focus of this work is on explaining experimental data using alternative models of learning in games. This work can be categorized according to the following typology. This is not an attempt at a thorough review of this literature, but a sketch with examples.

Myopic learning. One approach is to model agents as learning based only on their own payoffs to past strategy choices. This category is exemplified by the work on reinforcement learning. Agents are modeled as learning to play strategies that paid off for them in the past. The work of Roth and Erev (1998) is the best-known example, though others are developing reinforcement-based models and experiments (e.g., Slonim and Warwick, 1999). Agents are myopic in the sense that they do not look outside themselves when deciding which strategy to play, but rather pay attention only to their own payoffs from previous strategy choices. While models based on reinforcement learning generally "organize the data" well, the agents in these models do not learn fast enough. It can take many more trials of a simulated exchange to reach a convergence that occurs after just a few experimental trials with human learners.

Anonymous learning. These models assume that agents learn not just from their own experience, but also from observing the behavior of others. Other agents are anonymous in the sense that agents to not distinguish among them. This category incorporates most belief-based learning. Agents forecast the actions of others, based on the previous play of the agents. I also would include in this category quantal-response models (McKelvey and Palfrey, 1995, 1998; Anderson, Goeree, and Holt, 1997, 1998; Goeree and Holt, 1999), where agents forecast the distribution of actions and best-respond to that. These models tend to assume that beliefs are formed by some particular process. In a recent variation on this approach, Schotter and Nyarko

(1998) elicit beliefs, rather than assume a particular belief-formation rule. Shachat and Walker (1997) explore the rules that players use in learning, giving some insight into the prevalence of various degrees of sophistication in learning. Camerer and Ho (1998) develop a general model that subsumes both reinforcement learning and belief-based learning as special cases.

Social learning or learning with identity. In models with identity, agents differentiate others, and can learn differently from different persons. Several researchers have examined models where agents interact only with a subset of the population, their "neighbors". However other aspects of social structure have not received much attention. There is considerable evidence that people put different weights on the actions of others, depending on connections or status differences, expertise. They might copy who is doing well, or conform to what others are doing. It is this category that we wish to focus on.

While there has been a great deal of experimental research on learning, very few researchers have examined interactions with heterogeneous agents or hierarchical structure. Offermans and Sonnemans (1998) structure their experiment so that subjects observe the success of others. They find that people learn both from experience and by imitating successful others; subjects imitate the forecast of successful players when given the opportunity. Brown (1994) notes the importance of reference points in determining the path of the learning process; though he does not consider this case, such reference points could be provided by the observed decisions or advice of higher-status players.

Local interaction models of Ellison and others explore equilibria of systems with boundedly rational agents, but these models seldom explore the potential role of status differences. An exception is the work of Bala and Goyal (1998), who present a model of learning from neighbors where agents observe their close neighbors, but also observe "the royal family", a small set of agents observed by everyone. Because everyone sees them, they are unduly influential, and may cause the society to veer to a suboptimal equilibrium. In the context of technology adoption, they show that the structure of information flows can lock in an inferior technology.

III. Evidence for Status-Based Learning

Sociologists have long studied the connection between status and influence. (See Webster and Foschi, 1988, for a survey.) An early example is provided by Berger, et al. (1977), who argue that the degree of influence that one person has on another derives in part from status differences. They examine the effect of specific (such as task-relevant expertise) and diffuse (such as rank that is not task-relevant) status characteristics on influence in an experimental setting. Subjects are of high or low status, and face a counterpart with high or low status. Subjects make an initial choice in a binary-choice decision task, then they observe the other person's choice, then make a final choice. They find that both specific and diffuse status characteristics rank can significantly influence choices in a simple decision-making exercise. Both high and low-status subjects are more likely to imitate the choice of a high than a low-status counterpart's choice.

Few experiments in economics address the effect of status on decisions, and none investigate directly the issue of influence. Ball and Eckel (1996) award half of their subjects higher status, then match up players in an ultimatum game.[1] They find that while higher-status proposers tend to make higher offers, both higher and lower-status proposers offer more to higher-status respondents. Ball, Eckel, Grossman and Zame (1998) use the same procedure for

[1] In an ultimatum game two agents are provisionally allocated a fixed amount of money. One subject, the proposer, must propose a split of the pie. The other, the respondent, chooses to accept or reject the offer. If he accepts, the money is split as planned; if rejected, both players receive zero.

awarding status to show that relative status can affect prices in a box-design auction market.[2] On average, higher-status subjects earn 11.4 percent more than their lower-status counterparts.

Further evidence of the connection between status and influence is provided by the prevalence of celebrity endorsements. This strategy is a very popular choice for advertisers; more than ten percent of television advertising includes celebrity endorsement (Walker, et al, 1992). This form of advertising is popular because it is successful: Agarwal and Kamakura (1995) use event study methodology to show that announcements of celebrity endorsement contracts are associated with increases in stock value of firms.

Endorsement by a celebrity can be a bit risky: O.J. Simpson's fall from grace left advertisers worrying if their products were damaged by association. Hertz dropped Simpson in the wake of double murder charges; similarly Michael Jackson was dropped by PepsiCo after accusations that he molested a young boy. (Business Week, 7/4/94). On the other hand, Michael Jordan's return to the NBA had a positive impact on products he endorsed; Mathur, et al. (1997) estimate that his move resulted in an increase in value for his client companies of around 2 percent, or more than $1 billion in stock market value.

Celebrity advertising may be successful because humans are especially attuned to high-status individuals, and are more likely to learn from them (to adopt their advice) than from others. Cummins (1998) proposes "dominance theory" as a framework to account for human cognitive capabilities that cause them to be attuned to social norms involving status hierarchies. In related work, Paul Gilbert (1990) argues that emotional responses to status differences are

[2] In a box-design market, all sellers have the same cost and all buyers have the same reservation value. An equal number of units is demanded and supplied, leading to a vertical overlap in the two curves. They use an oral double auction institution, which pushes the market strongly to a uniform price. However, the equilibrium can be anywhere between the cost and reservation value. They find that in sessions where buyers have higher status, the market price is lower than in sessions where the sellers have higher status, leading to higher average earnings for high-status players.

strong motivators for behavior. Envy in particular motivates humans to imitate those who have what we want. Perhaps celebrity advertising evokes that desire to imitate.

In a different vein Gil-White and Henrich (1999) argue that status is a valuable signal that someone is worth paying attention to. Instead of focusing on dominance hierarchies, which are observed in many nonhuman as well as human societies, they distinguish between imposed status (dominance) and freely-conferred status (prestige). They argue that prestige is a second avenue to status and status-competition in humans that results from the combination of an intense group social life and the unique imitative capacities of humans. In a prestige hierarchy, selective attention to higher-status individuals enhances cultural transmission, which may lead to the reproduction of successful strategies. On an individual level, paying attention to those with prestige (and emulating them) can be beneficial, as people with status typically are successful in a variety of domains, and presumable have valuable knowledge. Selective attention to more successful persons leads to the transmission of more valuable information. Gil-White and Henrich argue that not only is this a reasonable learning shortcut, but it is an evolutionarily stable strategy. High status individuals have learned how to succeed in their environment, so imitating their behavior is reasonable. This argument helps make sense of the success of celebrity endorsement. If humans are attuned to high-prestige individuals, then endorsements by them can influence consumption and demand for particular products.

IV. Experimental Design

We designed experiments to test the influence of a high-status player in a simple coordination game. Two different games were used in the experiment, shown Table 1. These games replicate two of the coordination games (games 4 and 6) used in the study of equilibrium selection by Cooper, et al. (1990). The games have several interesting properties. First, there are two equilibria in each game, given by the strategy pairs (1,1) and (2,2). The second equilibrium carries higher payoffs for both players. However, in game 4, the equilibrium (1,1) is less risky, in the sense that if the player puts any prior probability on their opponent choosing the dominated strategy 3, choosing strategy 1 avoids the possibility of receiving zero. The choice of a strategy in these games is akin to the choice of a product by an individual, where one's choice is more valuable when others choose the same product.

Table 1 here

In their study, subjects are randomly rematched with other players and play one of the games in Table 1 for in each of 22 periods. Empirical results by Cooper et al. (1990) show that there is a tendency to choose action 1 in Game 4 (this is done 78% of the time in the last 11 periods). The tendency for subjects to choose action 2 is more pronounced in Game 6 (97% of the time in the last 11 periods). There is some learning that takes place in the early periods of their games, but by period 15 subjects almost always play action 1 in Game 4 and action 2 in Game 6.

We alter their design by adding a commonly-observed agent. This agent was a fictitious player; while in general we prefer not to use fictitious players, in this case experimental control was strongly enhanced by being able to specify the information about the strategy-choice of the commonly-observed agent. In future work we plan to introduce a procedure that selects a commonly-observed agent from among the subjects.

We deliberately tried to induce higher status for the fictitious player. Before beginning the experiment subjects were given a "generalized knowledge" quiz made up of 15 true/false questions (see Appendix). Each subject's quiz was graded, the subject was given the score and then the subject was told that Member C (the fictitious player) had answered all 15 items correctly and had completed the quiz faster then anyone else in the experiment. A score of 15 was credible as half the subjects correctly answered 10 or more items (with an average score of 10.04).

A total of 8 different manipulations were run with 15 different sessions. Subjects in each session played a single game for 30 periods, in two 15-period phases. A 2x2x3 incomplete factorial design was used in these experiments. The first factor pertained to the type of game used, as explained above. The second factor involved the order in which the stimulus was applied, whether in the first 15 periods or in the last 15 periods of the session. The final factor concerned the stimulus. In this experiment, when the stimulus condition was turned on, subjects were given information about the past move of the fictitious player. That player's moves were the same across 15 periods. In one condition, it was reported that the player had previously chosen row 1, in the second condition, row 2 and in the final condition, row 3. To control for the effect of this information, in the non-stimulus condition subjects were reminded of what their previous partner had selected. In all conditions subjects were also reminded of their own choice in the previous period.

The design is incomplete because several treatment combinations were ignored. In pilot tests it was clear that subjects easily ignored Member C's past behavior in some circumstances. If C was reported as routinely choosing strategy 3 in the second 15-period phase of the experiment, this probably was not very credible. By the second phase of the experiment, all of the pilot sessions had converged to an equilibrium at (1,1) or (2,2), and subjects had no incentive to leave those equilibria. Second, Game 6 converges so quickly to the (2,2) equilibrium on its own that we thought there would be little to learn by running sessions in which Member C's past behavior was reported as row 2.

V. Procedure

A total of 144 subjects were recruited from the student populations at Rice University (7 sessions) and Virginia Tech (8 sessions). Subjects were recruited from recruitment pools built by the authors and through email solicitations. The recruitment pools draw broadly from the student populations; participants were initially recruited from introductory social science courses. While some subjects previously had participated in decision making experiments, none were familiar with this particular design. The sex composition of subjects was almost evenly divided (52.8 percent were males) and subjects were mostly Caucasian (72.2 percent). When recruited, subjects were told they would be given a show-up fee of $5 and that they could earn additional cash during the course of the experiment. Subjects were not told how much they could earn, and if pressed, were told they could make considerably more than the minimum wage for less than 60 minutes of their time in the laboratory.

When subjects showed up at the laboratory they were randomly assigned to a computer. Subjects were seated at computer carrels that prevented them from seeing one another's screen or communicating with one another. All experiments conducted by a female experimenter who read a standard protocol, cautioning subjects not to speak with one another during the course of the experiment and to direct all questions to the experimenter. Subjects then proceeded to self-paced instructions given at their computer screen. The instructions were modified slightly from those given in the appendix for Cooper et al. (1990), and are reproduced in the appendix. At various places in the instruction subjects were tested for comprehension and corrected before being allowed to continue. In a post experimental questionnaire, 140 of 144 subjects agreed that the instructions were clear.

Subjects faced a 3x3 matrix and were told to pick an action. The matrix looked like a standard game matrix in normal form, with the subject's choices labeled as a row number and the counterpart's choices as column numbers. While Battalio et al. (1996?? -- van Hyck et al.?) note that population characteristics arise in these coordination games from labeling players as row and column players, we tried to avoid this by assigning all players as row players. The computer adjusted payoffs in each of the cells so that all players viewed themselves as row players and their counterparts as column players. Subjects were told that they would be randomly assigned a new counterpart in each period and that they might face the same person more than once. As few as 6 subjects and as many as 12 subjects were used in each experimental session; no significant difference was observed across sessions due to the number of players.

Subjects participated in 30 distinct periods and they were told this in advance, so it was clear that they might play one another multiple times. Players were all assigned an identity at the outset -- all were told that they were Player B. No mention was made of the identity of their counterpart in any period. Consequently, subjects had no basis for knowing with whom they were matched at any period.

Once subjects made a row choice, the computer told them to wait. Once everyone made their decision, then everyone was told the outcome for that period. Subjects were given a sheet of paper and a pencil and told to write down their earnings in each period. All earnings were given in experimental francs and subjects were told that the official rate of exchange was 90 francs to the dollar. Subjects were told they would participate in two phases of the experiment, with each phase lasting 15 periods. Subjects were told in advance that they would only be paid for one period in each phase of the experiment. At then end of a phase, subjects were presented with 15 electronic cards and asked to pick one. Once selected, the card was flipped over and displayed the period that was randomly chosen for payment. Subjects were told to circle that period, the computer translated the francs to dollars and subjects were told to continue with the experiment. At the conclusion of the experiment subjects were paid in cash and in private.

VI. Results

Results are shown in Figures 1-3 and Table 2. Figure 1 shows aggregate results by game and treatment. Recall that for game 4, (1,1) is the low-risk equilibrium, while (2,2) is payoff-dominant. For game 4, 63.2 percent play strategy 1 in the absence of a commonly-observed agent with 31.7 percent playing strategy 2. When subjects observe that Player C chooses 1, 77.9 percent play strategy 1, with only 17.1 percent choosing strategy 2. This difference is statistically significant (χ^2=38.57, df=2, p=.001), indicating that a Royal move enhances coordination on strategy 1. While average earnings are not significantly affected, the variance of earning is reduced in the sessions where Player C chooses 1.

A similar story can be told when subjects observe Player C choosing 2. Under that condition 52.2 percent play strategy 2, which increases the occurrence of the payoff-dominant equilibrium (2,2). When compared with the baseline condition the difference is again statistically significant (χ^2=68.96, df=2, p=.001). A sufficient number of subjects play strategy 2 to significantly increase earnings; compared with the baseline condition, in the sessions where Player C choses 2, average earnings increase by a statistically significant 7 percent (t=3.67, p=.0003).

When subjects are told that Player C chooses 3, subjects do not imitate that strategy. While (3,3) would give all subjects a higher payoff, it is not an equilibrium. Instead of imitating C, subjects best-respond by choosing strategy 1, and they do so at rates higher than under the baseline condition (χ^2=44.51, df=2, p=.001). The effect is to reduce average earnings slightly, and to increase the variance of earnings. Observing a Royal move leads subjects toward one equilibrium or another. However, it does not induce subjects to play a very risky, non-equilibrium strategy.

The results for Game 6 tell a different story. Recall that for Game 6, (2,2) constitutes a very powerful equilibrium. Here neither message has a much weaker impact on subjects' strategy choice. Under the baseline condition subjects are playing strategy 2 more than 90 percent of the time. When subjects are told that Player C chose strategy 1, the number of subjects who play that strategy doubles, rising from 7.5 percent to about 15 percent of the plays (the difference is statistically different, χ^2=29.99, df=2, p=.001). When told that Player C chose 3. their strategy is affected in the same way as Game 4. Subjects choose 1, the best response to strategy 3, at a rate that is triple the baseline rate, and the difference is statistically significant (χ^2=55.27, df=2, p=.001). In both cases, announcing a move by royalty has a small but significant negative impact on earnings. (For royalty = 1, t=4.25, p=.0001; for royalty = 3, t=6.93, p<.0001). What is interesting is that a Royal move can shift some players toward an inefficient equilibrium (although this effect is small). The same cannot be said about a non-equilibrium move that is strictly dominated.

Results by treatment are shown in Figures 2 and 3. These figures illustrate the changes in the distribution of strategies as the games progress across 30 periods of play. Note in Figure 2 that announcing to subjects Player C's choice of 2 has a strong apparent effect on Game 4, whether it comes in Phase 1 or Phase 2. The effect of move 3 by Player C is small and has nearly disappeared by the end of Phase 1. In Figure 3, it appears that announcing that Player C chose 1 has some effect on Phase 1 play, but none if it occurs in Phase 2. If subjects are told that

Player C chose 3, some respond by playing 1; however, this behavior dies out by the end of phase 1.

Table 2a-b contains the results of regression analysis of the data. We estimate a multinomial logit model to estimate the effect of the commonly-observed agent on the subjects' choices among the three available strategies, controlling for the individual agents' history. Because the strategy choices are not ordered, we estimate each model using a multinomial logit for unordered response variables. In each instance the dependent variable of interest is the strategic choice by a player.

In model I, shown in Table 2a, we model the time structure of the game with three variables in addition to the intercept. Recall that each experiment consisted of two fifteen-period phases. From Figures 2 and 3, it is clear that subjects learn over time which strategy to play, and tend to converge to one of the strategies associated with an equilibrium. Our first variable, then, is a simple count for the period in which the subject is making a decision. To model the data as a pooled, interrupted time series, we introduce a second variable set to 0 for the first 15 periods and then counting from 1 for periods 16 through 30. This produces an estimate of a new slope for the second phase of the experiment.[3] To accommodate a shift in intercept with the phase-two regime, we also include a dummy variable that is 0 for the first 15 period and 1 thereafter. Finally, we include a series of dummy variable standing for each of the experimental conditions. For example, if the fictitious player, Member C, reported moving 2 in the previous period, then the associated dummy variable took on a value of 1 and 0 otherwise.

The multinomial logit estimates are reported with respect to a default strategy. Since in the baseline treatments Game 4 converged to an equilibrium at (1,1) and Game 6 to (2,2), we use

[3] This specification imposes a linear structure on the subjects' learning. Including the log of the period number to model nonlinearity did not change the estimates.

these as the default outcomes. Estimates for Games 4 and 6 are reported separately. The coefficients should be interpreted as contributing to the likelihood of playing a specific strategy relative to the default. In Table 2 the first column for game 4 represents the likelihood that a subject would play strategy 2. The intercept is strongly negative, indicating that the average choice was strategy 1 (which of course can be seen from the charts in figure 2). Interestingly, there is almost no effect for learning over time. Each of the dummy variables representing the experimental conditions are significant. If a subject sees a royal move of 1, then the subject is less likely to choose strategy 2. Similarly, if subjects see a royal move of 3 they are also less likely to choose 2. It is when subjects see a royal move of 2 that they are increasingly likely to pick strategy 2.

Column 2 reports the estimates for strategy 3. Here the intercept is negative and the overall time trend for the experiment is negative. Subjects are less likely to play strategy 3 at the outset, and the popularity of 3 diminishes over time. There is a strong intercept effect at the second phase, consistent with a "restart" effect observed in many experiments. Here subjects return to "experimenting" with strategies to see if others will shift their play. However, this quickly returns to previous patterns of play. One the dummy variables for the experimental conditions carries a significant coefficient; if subjects observe a royal move of 1, they are less likely to pick strategy 3 (and this is marginally significant). In sum, subjects are responding to move by Member C when they make sense. In game 4 this means playing either strategy 1 or strategy 2.

<Table 2 About Here>

In game 6, the risk associated with playing 2 is removed, and the joint play of strategy 2 leads to an equilibrium that is Pareto-dominant. The third column of Table 2 represents the likelihood of choosing strategy 1 given the reference point of strategy 2. Here the coefficient on the period number indicates that subjects are less likely to choose strategy 1 over time. There is no additional learning in the second phase of the experiment -- and little wonder because by that point subjects have settled into the (2,2) equilibrium. There is a strong, positive effect when subjects observed a royal move of 1. The same is true when they observe a royal move of 3 -- although here it is clear that subjects are not following but instead playing best response against strategy 3. The fourth column provides estimates for the strategy choice of 3. The intercept and time variable carry significant negative coefficients, interpreted as in Game 4. Subjects again exhibit a restart effect at the beginning of phase 2. In effect a few individuals are playing strategy 3 in the first several periods of the second phase, but quickly return to playing strategy 2.

Table 2b contains results for a model with several new variables added to capture individual histories. We estimate a one-period adaptive learning model, and include the subject's own past strategy and his counterpart's strategy in the previous round as variables. The columns of Table 2b are the same as Table 2a. In general, the time structure variables show the same pattern as the previous estimates. Again the coefficients on the experimental condition variables are all in the expected direction, and statistically significant. If the royal move is 2, subjects are more likely to play strategy 2; a royal move of 1 or 3 both reduce the play of 2. An individual's prior behavior is a good predictor of future play-- if the individual has previously chosen strategy 1 then she is likely to again choose 1. This is partly offset by the experience of having a partner play 2, which has a positive effect on choosing strategy 2. By comparison, there is very little that is correlated with a player choosing strategy 3 -- except the fact that the individual had previously played 3. However, over time a member is decreasingly likely to play strategy 3.

In game 6 the best predictor of whether a subject will pick strategy 1 is whether that individual played strategy 1 in the preceding move. This tendency is somewhat offset by a partner choosing strategy 2 in the previous period. It is also notable that the effects of the experimental conditions diminish to insignificance in game 6. This is especially true for strategy 3. A subject's past play of strategy 3 is associated with a greater likelihood of playing 3, and a prior partner choosing strategy 2 reduces the play of 3. However, apart from individual histories, observing a royal move has no effect. This supports our earlier point that a the equilibrium at (2,2) is especially strong.

VII. Discussion and Conclusion:

The results of our experiments indicate that the presence of a commonly-observed agent can influence subjects to coordinate on an attractive strategy more quickly, and can move subjects to a higher-payoff strategy even when it is risky. However, this common information cannot move subjects away from an equilibrium that dominates in terms of payoffs and risk, and it cannot move a population to a strategy that is not an equilibrium.

These results are consistent with the notion that people learn differently from different people. Observation of the strategies of others may not always lead to the same response or degree of imitation. Some people have more credibility than others, so their behavior is worth paying more attention to. A cheap rule of thumb may be to copy high-status people. As with celebrity endorsements, the observed choice of a higher-status decision-maker may have a strong influence on what others do. This influence appears to be strongest when the choice is among two or more "good" outcomes, but not when one of the choices is clearly inferior. This might lead us to expect more celebrity endorsements when products are very similar but not easily ranked on the basis of quality, such as soft drinks, non-prescription medicines, etc. It is important to investigate further the connection between status and influence.

Several researchers have argued that solving games with multiple equilibria requires that one of the equilibria be salient in some way. Most of the research on focal points in decision-making derives from Schelling's 1960 book. In our experiments, by announcing Player C's move, we in effect make the outcome associated with that move more salient or focal. If equilibrium selection were based solely on focal points and our messages made a given strategy focal, the message should be equally powerful for all games. However, we see that the impact of the strategy announcement depends on the game context, and also on whether the observed strategy is sensible. This reduces the likelihood that focal points alone are responsible for the influence of the commonly-observed agent.

Table 1: Game Matrices
All payoffs are given in experimental francs; 1 franc = $.09

Game 4.

	Col. 1	**Col. 2**	**Col. 3**
Row 1	350,350	350,250	700,0
Row 2	250,350	550,550	0,0
Row 3	0,700	0,0	600,600

Game 6.

	Col. 1	**Col. 2**	**Col. 3**
Row 1	350,350	350,250	700,0
Row 2	250,350	550,550	650,0
Row 3	0,700	0,650	600,600

Table 2a
Coefficients, Standard Errors in Parentheses, Significance Levels in Italics

	Game 4		**Game 6**	
	Strategy 2| *Strategy 1*	*Strategy 3|* *Strategy 1*	*Strategy 1|* *Strategy 2*	*Strategy 3|* *Strategy 2*
Intercept	**-1.034** (.158) *p=6.557*	**-.979** (.208) *p<.001*	-.203 (.195) *p=.299*	**-.750** (.291) *p=.010*
Time (1 … 30)	.0109 (.018) *p=.534*	**-.167** (.031) *p<.001*	**-.180** (.025) *p<.001*	**-.300** (.051) *p<.001*
Time 2nd Phase (1 … 15)	.015 (.022) *p=.498*	.048 (.049) *p=.324*	.042 .046 *p=.360*	**.162** (.081) *p=.046*
Phase 2 Dummy	.078 (.189) *p=.678*	**1.296** (.413) *p=.002*	.541 (.352) *p=.124*	**1.926** (.702) *p=.006*
Member C Moved 1	**-.765** (.142) *p<.001*	**-.471** (.266) *p=.077*	**.482** (.186) *p=.010*	.510 (.301) *p=.090*
Member C Moved 2	**.989** (.110) *p<.001*	.041 (.273) *p=.881*	--	--
Member C Moved 3	**-.895** (.199) *p<.001*	-.264 (.278) *p=.342*	.407 (.212) p=.055	-.342 (.430) *p=.427*
	LLF=-1967.47	r^2=0.073	LLF=-759.81	r^2=0.153

Table 2b
Coefficients, Standard Errors in Parentheses, Significance Levels in Italics

	Game 4		**Game 6**	
	Strategy 2| Strategy 1	*Strategy 3| Strategy 1*	*Strategy 1| Strategy 2*	*Strategy 3| Strategy 2*
Intercept	**-.715** (.251) *p=.004*	**-.979** (.291) *p=.001*	-.332 (.304) *p=.275*	**-.818** (.373) *p=.029*
Time (1 … 30)	.022 (.022) *p=.324*	**-.119** (.037) *p=.001*	**-.110** (.030) *p<.001*	**-.212** (.061) *p=.001*
Time 2nd Phase (1 … 15)	-.021 (.027) *p=.442*	-.001 (.054) *p=.984*	-.002 (.049) *p=.963*	.087 (.088) *p=.327*
Phase 2 Dummy	-.019 (.219) p=.931	**1.141** (.436) *p=.009*	.337 (.372) *p=.362*	**1.513** (.729) *p=.038*
Member C Moved 1	**-.387** (.163) p=.017	-.238 (.285) *p=.403*	.201 (.207) *p=.331*	.146 (.354) *p=.681*
Member C Moved 2	**.557** (.133) *p<.001*	-.043 (.291) *p=.882*	--	--
Member C Moved 3	**-.499** (.221) *p=.024*	-.039 (.307) *p=.898*	.176 (.236) *p=.457*	-.315 (.471) *p=.504*
Previous Move=1	**-2.083** (.414) *p<.001*	-.917 (.485) *p=.059*	**1.206** (.507) *p=.017*	.470 (.753) *p=.533*
Previous Move=2	.224 (.413) *p=.587*	.621 (.498) *p=.212*	-.157 (.498) *p=.753*	.268 (.671) *p=.690*
Previous Move=3	**-.938** (.457) *p=.040*	**1.702** (.494) *p=.001*	.644 (.577) *p=.264*	**2.247** (.704) *p=.001*
Prior Partner's Move=1	.500 (.293) *p=.088*	-.472 (.342) *p=.168*	.178 (.369) *p=.629*	-.801 (.542) *p=.139*
Prior Partner's Move=2	**1.667** (.299) *p<.001*	.043 (.377) *p=.910*	**-.881** (.356) *p=.013*	**-1.220** (.484) p=.012
	LLF=-1621.68	r^2=0.236	LLF=-709.27	r^2=.210

Figure 1: Aggregate Results by Game and Treatment

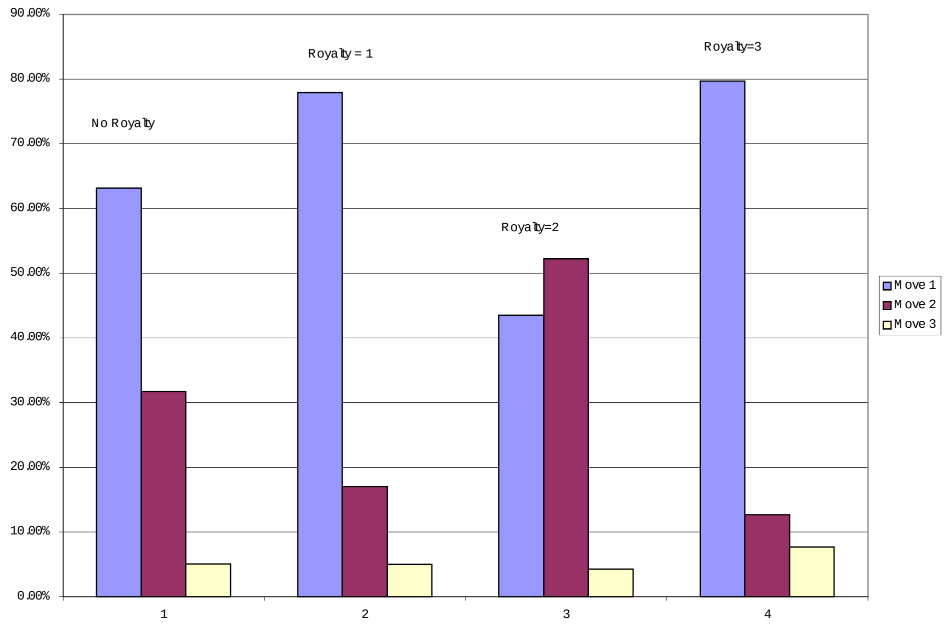

a. Game 4

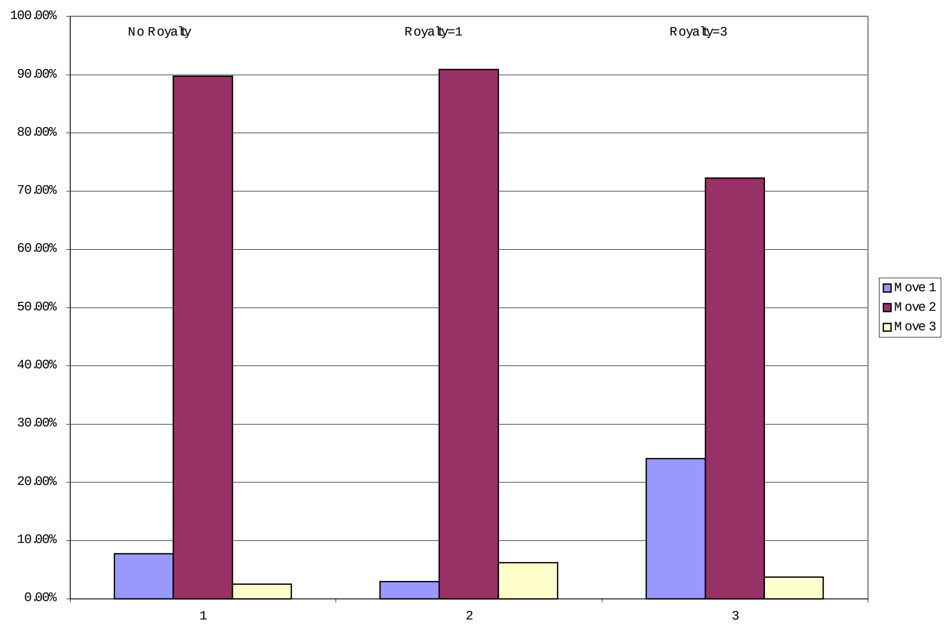

a. Game 6

Figure 2: Results by Treatment, Game 4

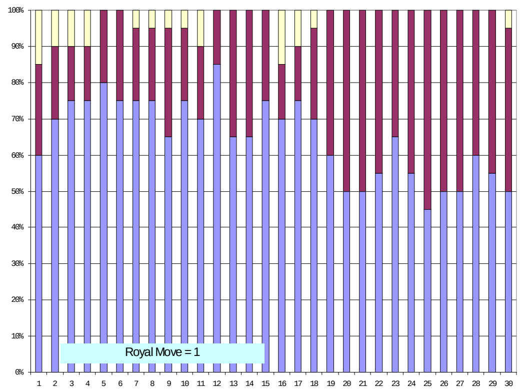

Royal Move = 1, Phase 1

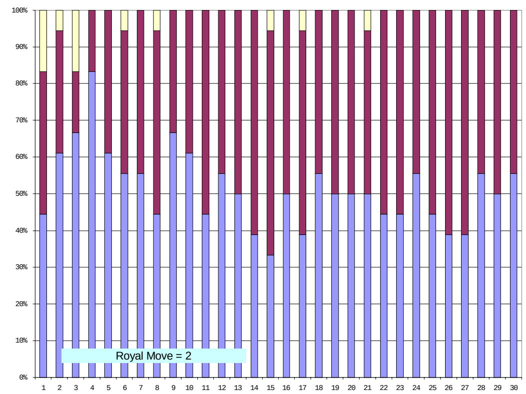

Royal Move = 2, Phase 1

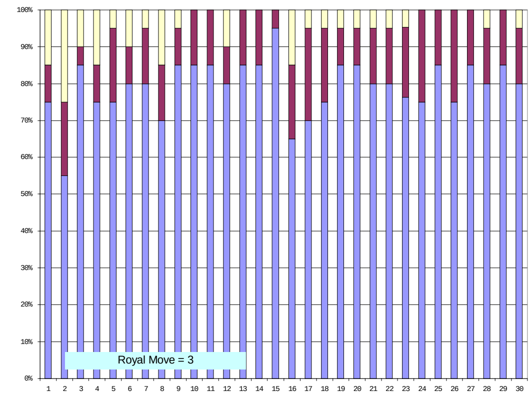

Royal Move = 3, Phase 1

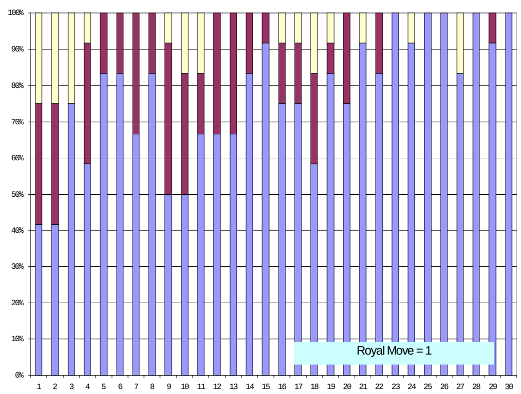

Royal Move = 1, Phase 2

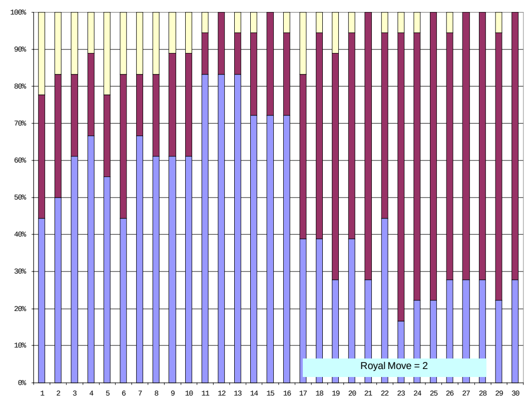

Royal Move = 2, Phase 2

Figure 3: Results by Treatment, Game 6

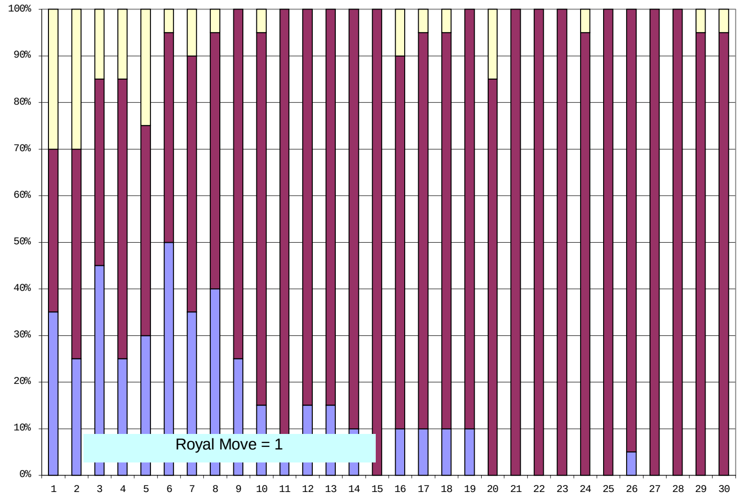

Royal Move =1, Phase 1

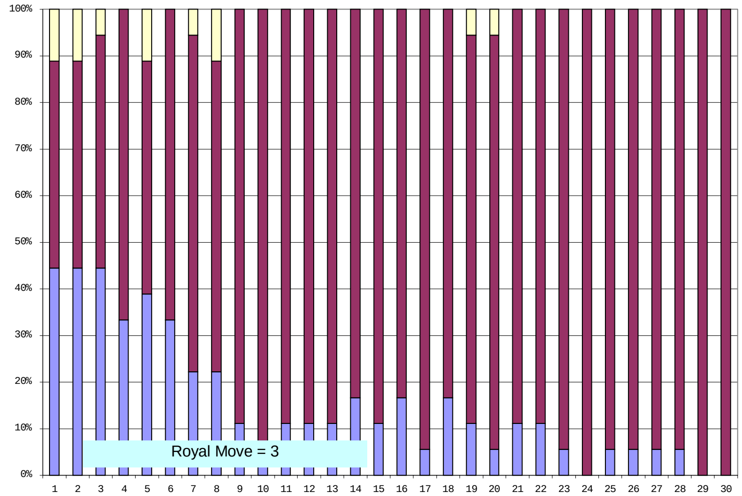

Royal Move =3, Phase 1

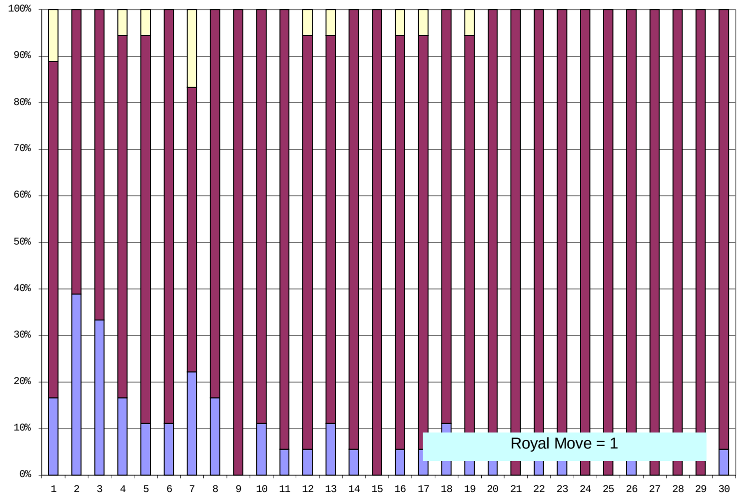

Royal Move =1, Phase 2

Social Learning: References

Agrawal, Jagdish, and Wagner A. Kamakura. 1995. "The Economic Worth of Celebrity Endorsers: An Event Study Analysis." *Journal of Marketing* 59: 56-62. July.

Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1997) "Stochastic Game Theory: Adjustment and Equilibrium with Bounded Rationality," University of Virginia, Unpublished Manuscript.

Anderson, Simon P., Jacob Goeree, and Charles A. Holt (1998) "Control Costs and Equilibria in Games with Bounded Rationality," University of Virginia, Unpublished Manuscript.

Armstrong, Larry. 1994. "Still Starstruck." *Business Week*, July 4, 1994, p. 38.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbors." *Review of Economic Studies* 65: 595-621.

Ball, Sheryl, and Catherine Eckel, "Buying Status: Experimental Evidence on Status in Negotiation." *Psychology and Marketing*, 13(4):381-405. 1996.

Ball, Catherine Eckel, Philip Grossman, and William Zame. 1998. "Status in Markets." Unpublished manuscript, Department of Eocnomics, Virginia Polytechnic Institute and State University.

Joseph Berger, M. Hamit Fisek, Robert Z. Norman, and Morris Zelditch, Jr., *Status Characteristics and Social Interaction: An Expectation-States Approach*. New York: Elsevier. 1977.

Brown, Paul M.; 1994, "Learning from Experience, Reference Points, and Decision Costs", *Journal of Economic Behavior and Organization* 27, 381-399.

Camerer, Colin, and Teck-Hua Ho. 1998. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica* 67:827-874

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, Thomas W. Ross. 1990 "Selection Criteria in Coordination Games: Some Experimental Results." *American Economic Review* 80(1):218-233.

Cummins, Denise D. 1998. "Social Norms and Other Minds: The Evolutionary Roots of Higher Cognition." In *The Evolution of Mind*, edited by D. D. Cummins and C. Allen. New York: Oxford University Press, p. 30-50.

Gilbert, Paul. 1990. "Changes: Rank, Status and Mood." In *On the Move: The Psychology of Change and Transition*, edited by S. Fischer and C. L. Cooper, pp. 33-52. New York; Wiley.

Gil-White, Francisco and Joe Henrich. 1999. "The Evolution of Prestige: Freely-conferred status as a mechanism for enhancing the benefits of cultural transmission." Unpublished paper. UCLA Department of Anthropology.

Goeree, Jacob K., and Charles A. Holt (1999) "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96(September), 10564-10567.

Mathur, Lunette Knowles, Ike Mathur, and Nanda Rangan. 1997. "The Wealth Effects Associated with a Celebrity Endorser: The Michael Jordan Phenomenon." *Journal of Advertising Research* 37: 67-73

McKelvey, Richard D., and Thomas R. Palfrey. 1995. "Quantal Response Equilibria in Normal Form Games." *Games and Economic Behavior* 7: 6-38.

McKelvey, Richard D., and Thomas R. Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1(1): 9-42.

Offerman, Theo; and Sonnemans, Joep; 1998, “Learning by Experience and Learning by Imitating Successful Others,” *Journal of Economic Behavior and Organization* 34, 559-575.

Roth, Alvin E. and Ido Erev (1998) "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed-Strategy Equilibria." *American Economic Review* 88(4): 848-881. September.

Schelling, Thomas, 1960. *The Strategy of Conflict*. Cambridge: Harvard University Press.

Nyarko, Yaw, and Andrew Schotter (1998) “Equilibria in Beliefs and Our Belief in Equilibria,” New York University, Discussion Paper, presented at the Summer 1998 ESA

Meeting.Slonim and Warwick. Paper presented at Penn State conference.

Van Huyck, John B., Battalio, Raymond C., and Rankin, Frederick W. 1997. On the Origin of Convention: Evidence from Coordination Games," *Economic Journal*; (107) 442:576-96.

Walker, Mary, Lynn Langmeyer, Daniel Langmeyer. 1992. "Celebrity Endorsers: Do you get what you pay for?" *Journal of Consumer Marketing* 9: 69-76. Spring.

Webster, Murray, Jr. and Martha Foschi. 1988. "Overview of status generalization." Chapter 1 in Murray Webster, Jr. and Martha Foschi (Eds.), *Status Generalization: New Theory and Research*. Stanford: Stanford U. Press.