The Effect of Information in Learning to Play Nash

Chris Geissler

May 9, 2006

Abstract

I design an experiment that focuses on the role that information plays when agents are choosing a strategy. I find that giving players more information about the strategies and payoffs that other participants are using may make the game more or less likely to reach the point where players are playing at the Nash equilibrium, depending on characteristics of the payoff functions. In one form, where the players' payoffs are strategic complements, providing subjects with more information decreases the chance that they will play at the Nash equilibrium. In the other form, where their payoffs are mixed, increasing the level of information makes players more likely to play at the Nash and best respond to one another.

When making decisions, an agents uses a number of resources. Among the most important are her past experiences as they often provide some intuition for which strategies are best. In addition, a subject is likely to look at what strategies other agents are employing if possible to see if they are worth imitating. An individual's ability to learn about which strategies are effective within a game is dependent on the amount of information that she receives. One would expect that if a subject is given more information about the structure of a game's payoffs and the actions other players are choosing, she will figure out the strongest strategies more quickly. It therefore seems likely that games will move towards their Nash equilibria faster when subjects are given more information.

There have been a number of studies that have looked at how human subjects learn when playing a specific repeated game. Harrison and McCabe (1996), Duffy and Feltovich (1999), and Bohnet and Zeckhauser (2004) all look at how learning affects subjects' choices in the ultimatum game. The studies had inconsistent results as Harrison and McCabe (1996) found that increasing the amount of information move the game towards Nash equilibrium play whereas both Duffy and Feltovich (1999) and Bohnet and Zeckhauser (2004) found that increasing the level of information made the game move away from Nash play. Erev and Rapoport (1998) analyze the effects of varied information in a market entry game and Stahl (2000) focuses on the the role of information in bimatrix games. Erev and Rapoport (1998) find that giving subjects information about other players' payoffs moves the game away from the Nash equilibrium and Stahl (2000) finds that subjects do use the information provided to learn as the game as played. The Cournot model has also been studied by Bosch-Domenech and Vriend (2003) and Huck, Normann and Oechssler (1999) as both gave subjects differing amounts of information about their payoffs. These studies found that when a subject was given more information about other players' outputs and payoffs, he typically chose a higher output thus moving away from the Nash equilibrium. When he received more information about his payoff, however, he was more likely to choose a lower output thus moving the game towards the Nash.

There is also an abundance of theoretical work which attempts to model how subjects learn. This includes Camerer (2003) which builds a model that draws from a number of learning models. While all of these papers look at the role of information in a subject's abilities to learn the best strategies of a game, they focus primarily on the role of information related to the payoffs, not other players' strategies. In the real world, actors often gain information by observing the strategies employed by those around them, suggesting that the role of information players receive about other players' strategies merits study. Only Bosch-Domenech and Vriend (2003), Erev and Rapoport (1998), and Duffy and Feltovich (1999) look at how varying the amount of information given to subject about the other players' strategies and payoffs affects the overall game and subjects' ability to learn. These studies only focus on a limited number of games and do not analyze games more complex than the Cournot oligopoly. This means that they do not effectively model how subjects reach more complicated decisions in some circumstances. In addition, none of these papers focuses on how this role of additional information may vary based upon certain traits in each subject's profit function.

The structure of the payoffs is important in shaping a subject's ability to learn and the probability that play will approach the Nash. An important characteristic of some profit functions is supermodularity. A game with continuous profit and best response functions is supermodular if for each player, the best response increases as his opponent's choice increases. There are numerous studies focusing on how supermodularity affects the probability that the game moves to the Nash equilibrium. Chen and Gazzale (2004) focus on compensation mechanism games and find that supermodular games are more likely to converge to the Nash than games which are not supermodular.

This study designs an experiment that focuses on the distinction between markets where players are only able to observe their own actions and those where they are also able to observe the decisions made by another subject. In addition, it uses multiple payoff structures to see if this information affects different games in a similar manner thus giving a two by two experimental design. I am interested in how each of these treatments affects the likelihood that a subject's strategies will approach his Nash equilibrium value. Additionally, I will test whether this additional information improves a subject's ability to best respond as this serves as another indicator of the value of this additional information. This may be an even better indicator of the importance of this information as a player may choose a Nash equilibrium value when it is not a best response. Since best responding maximizes a player's payoff, it is in a player's interest only to play a Nash equilibrium value if her opponent does as well whereas it is always in her interest to best respond regardless of her opponent's strategy.

The experiment looks at the differences between how individuals use information acquired by playing a strategy and information received by seeing the strategies of those around them. There are treatments with two different information levels. In the Low Information treatment, each players sees only his chosen value and payoff, and the choice made by the player he is matched with. In the High Information treatment, however, he is also given the value chosen and payoff of another randomly selected player of his type, which I call the public player. In addition, he sees the choice of the player matched with the public player. There are different ways of interpreting the role of this additional information. For example, one possibility is that a player who has played the Low Information game for two rounds, and thus has seen two sets of payoffs, has the same amount of information as a player who has played the High Information game for one round, and has again seen two payoffs, one of which is his own. If this were the case, the player in the High Information game would have twice as much information as the subject in the Low Information treatment in a given round.

There may be some fundamental differences between the information acquired through playing a game oneself and that learned by observing other players. By seeing others, a player may learn about strategies that he was not willing to try because he deemed them too risky or believed the expected profits were low compared with other strategies. Thus, by observing a number of other players over many rounds, he will likely get a broader perspective of various strategies as others may try these strategies. At the same time, some players might put a greater emphasis on information they acquired firsthand than that which they observed from other players of their type.[1]

It is also possible that a player initially chooses values which are similar to those he is observing. But over time, she notices that other players are playing lower values and getting higher payoffs than before whereas her payoff is not as high now as it was earlier in the game. She is therefore likely to change her strategy and play a lower value as well thus getting a higher payoff.

In addition to having two levels of information, Low Information and High Information, the experiment also uses two different sets of payoff functions. In each form, subjects are participating in a two player game. In the first, Strategic Complements, the value of a player's best response increases as her opponent's choice increases for both types of players meaning it is supermodular. In the second, the Mixed game, this trait holds for type-1 players, who have the same payoff function in each game, but not for type-2 players. In the Mixed game, the value of a type-2 player's best response decreases as the type-1 player's value increases meaning it is not supermodular. Assuming that each player best responds to his opponent's choice in the previous

[1] This concept is captured in Camerer (2003) with the inclusion of "imagination weight" into his learning model. Imagination weight discounts the weight an agent puts on a a strategy that he did not experience first hand. In Camerer's model, however, the subject does not see other player's strategies or payoffs so this discounting may be necessary because the subject has to calculate what profits they would have earned had they chosen a different strategy.

round, the equilibrium paths in these games are quite different. While it takes the same number of steps to reach the Nash in each, the path is more volatile in the Mixed game. This result is demonstrated in Figure 2 in the following section. This leads me to speculate that human subjects will likely have more trouble getting to the Nash in the Mixed game. This expectation is supported by Chen and Gazzale (2004) which found Nash play more likely in games where the profit functions are supermodular.

In many cases, the data is consistent with economic intuition and previous studies. In the Mixed game, the High Information treatment is more likely to converge towards the Nash equilibrium than the Low Information treatment. There are also some results that are not consistent with conventional economic theory. In the Strategic Complements game, players in the High Information treatment are less likely to best respond and the game does not converge towards the Nash equilibrium with the same probability as players in the Low Information treatment. In addition, subjects in the Strategic Complements High Information treatment are less likely to play near the Nash equilibrium than those in the Mixed High Information treatment.

1 General Economic Environment

The profit functions for both the Strategic Complements and Mixed treatments of the experiment are of the form given by

$$\pi_i(x_i, x_j) = ax_i + bx_j - cx_i^2 - dx_2^2 + ex_ix_j. \quad (1)$$

Player i's choice is represented by x_i. Her opponent's choice is x_j. This general profit function was chosen because even though subjects receive their payoff function, it is too difficult for most players to solve their best response function or find the Nash equilibrium as this would require doing multivariable calculus during the game. Yet a player is likely to develop an intuition for what strategies might be better or worse depending on what his opponent plays. Since this experiment is primarily concerned with the relationship between learning in games, subjects ability to best respond, and convergence to the Nash equilibrium, it is important that they do not know their best response function before the game begins. It is also necessary, however, that they have some understanding of the profit function as this prevents players from choosing values randomly throughout the game. This profit function allows many of the subjects to develop a general intuition for what their best response function is.

Solving for player i's best response function, I find that

$$BR(x_i) = \frac{a + ex_j}{2c}. \tag{2}$$

Plugging one best response function into the other, I find that the Nash equilibrium occurs at

$$\{x_1; x_2\} = \left\{ \frac{4c_1c_2}{4c_1c_2 - e_1e_2} * \frac{a_1 + \frac{e_1a_2}{2c_2}}{2c_1}; \frac{4c_1c_2}{4c_1c_2 - e_1e_2} * \frac{a_2 + \frac{e_2a_1}{2c_1}}{2c_2} \right\} \tag{3}$$

A game has strategic complementarities if, when strategies are ordered, a higher choice by a subject's opponent provides an incentive for him to also play a higher value. Supermodular games, defined as games where the incremental return to any player from increasing her strategy is a nondecreasing function of the strategy choices of other players, always have strategic complementarities Chen and Gazzale (2004). If each players' profit and best response functions are differentiable in the strategy space, then a game is supermodular if each player's marginal utility of increasing her choice rises as her opponent's value increases.

Several studies including Milgrom and Roberts (1991) and Milgrom and Shannon (1994) provide theoretical explanations which predict that supermodular games are likely to converge to the Nash equilibrium. Chen and Gazzale (2004) provide support for this theory and their laboratory results suggest that supermodular games typically provide easier learning environments for the subjects as the game they used had a higher convergence rate to the Nash equilibrium when it was supermodular than when it was not.

It is possible for two games to share several traits including the same Nash equilibrium and equal absolute values of slopes of the best response functions with only one of them being supermodular. A Nash equilibrium occurs in a two player game where their best response functions intersect. This means that at this point, neither player has an incentive to change his strategy if he assumes that the other player will also keep her strategy constant. If we hold one of these values constant, like with type-1 players above, both positively and negatively sloped lines can intersect this linear function in the same point if the negatively sloped line has a larger y-intercept. This is shown in Figure 1. Under such circumstances, both sets of functions will have the same Nash equilibrium.

In the following section, I introduce the specific profit functions and solve for each players' best response as well as the the Nash equilibrium.

2 Experimental Design

The experiment uses a two by two experimental design and focuses on two key questions that address learning in games. The first condition of the game that I vary is the profit function. In each of the two forms of the game, there is a different profit function for type-2 players giving the games different properties. In the first form, Strategic Complements, the supermodularity condition outlined in the previous section holds. In the second form, Mixed, the supermodularity condition does not hold. In addition, the experiment also gives the subjects different levels of information regarding the values that other players of their type are choosing and what payoffs they are receiving. In one treatment, Low Information, subjects only receive information about their payoff. In the High Information treatment, however, they are also given information about a randomly selected subject of their type, the public player, at the end of the round.

	Low Information	High Information
Strategic Complements	4 Sessions	4 Sessions
Mixed	4 Sessions	4 Sessions

Table 1: Types of sessions

Table 1 outlines the 4 treatments and the number of sessions of each that was run. The experimental procedures are described in more detail below.

2.1 *The Economic Environment*

In the following equations, let x_1 be the choice of a type-1 player and x_2 be the choice of a type-2 player. In all treatments, the payoff function for type-1 players is

$$\pi_1 = 204x_1 + 20x_2 - 18x_1^2 - 15x_2^2 + 30x_1x_2. \tag{4}$$

In the Strategic Complements treatment, a type-2 player's payoff function is

$$\pi_2 = 160x_2 + 60x_1 - 20x_2^2 - 13x_1^2 + 30x_1x_2.$$

In the Mixed game, actions are strategic complements for type-1 players, but are strategic substitutes for type-2 players. Equation 4 determines type-1 player profits, while type-2 players profits are

$$\pi_2 = 1280x_2 - 360x_1 - 16x_2^2 + 8x_1^2 - 24x_1x_2. \tag{5}$$

When substituted back into Equation 3, the subgame-perfect Nash equilibrium is (x_1^*, x_2^*)=(24,22) in both games. The payoff functions in each form of the game were chosen for several reasons. This Nash equilibrium does not lie the center of the strategy space[2] for either player therefore avoiding equilibrium convergence as a result of focal points.[3] By plugging these values

[2]The strategy space, S, is defined as the set of choices that a subject has when selecting his strategy. In this experiment, $S = (0,1,\ldots,40)$.

[3]A focal point, also known as a Schelling point, occurs at a strategy that subjects are likely to play without communication because it seems special. In this game, the most obvious focal point occurs at 20, the center of the strategy space.

into Equation 2, one can generate each player's best response function. In both games, type-1 players have the best response function given by

$$BR(x_1) = \frac{204 + 30x_2}{36}. \tag{6}$$

In the Strategic Complements game, type-2 players have the best response function given by

$$BR(x_2) = \frac{160 + 30x_1}{40}. \tag{7}$$

In the Mixed game, type-2 players have the best response function given by

$$BR(x_2) = \frac{1280 - 24x_1}{32}. \tag{8}$$

As these equations show, in each game, the profit functions and best response functions are differentiable for all values in the strategy space. Both type-1 and type-2 players have positive sloping best response functions for the Strategic Complements game but type-2 players have a negatively sloped best response function for the Mixed game. This can be easily shown as type-1 players always have a best response function with a slope of $5/6$. In the Strategic Complements treatment, a type-2 player's best response function has a slope of $3/4$. As a result, in Strategic Complements, a player's marginal utility of his increasing choice rises as his opponent's value increases for both player types. This condition does not hold for type-2 players in the Mixed game however. Because both types of players have positively sloped best response functions in the Strategic Complements game, it is supermodular. In the Mixed game, however, type-2 players have a best response function with a $-3/4$ slope meaning it is not a supermodular game.

In the experiment, each player chooses $x_i \in (0,1,\ldots,40)$. With the chosen profit functions, each game offers the possibility of players of either type earning significant losses in a particular round. When both players choose their respective Nash values, type-1 players earn 3548. In the Strategic Complements form, type-2 players earn 3632 at the Nash whereas they earn 3712 at the Mixed game Nash equilibrium.

Because equilibrium profits for each player are close to each other, the game is relatively fair to both types of players.[4] Also, even though the best response function for type-2 players varies depending on what form of the game is being played, the absolute value of the slope of the line remains the same. Therefore if a type-2 player eliminates possible strategies that are dominated[5], he will be able to eliminate the same number of strategies in each form of the game.

I present the best-response function in the Strategic Complements game in Figure 1(a), and those in the Mixed game in Figure 1(b). In the Strategic Complements game, they each have a positive slope of a different magnitude. If both players start by choosing values above the Nash equilibrium and Cournot best-respond[6] in the subsequent rounds, the values they choose will drop in each of the following rounds until they approach the Nash equilibrium

[4] This may be important since some studies, such as Duffy and Feltovich (1999) found that providing subjects with more information made it likely that they would play a more equitable strategy even if it gave them a lower monetary payoff. I expect that the relatively equal payoffs at the Nash equilibrium will prevent such strategies.

[5] A strategy is dominated if it is not a best response to any play from one's opponent.

[6] A subject Cournot best-responds when they play the value that is a best response to their opponent's strategy in the previous round.

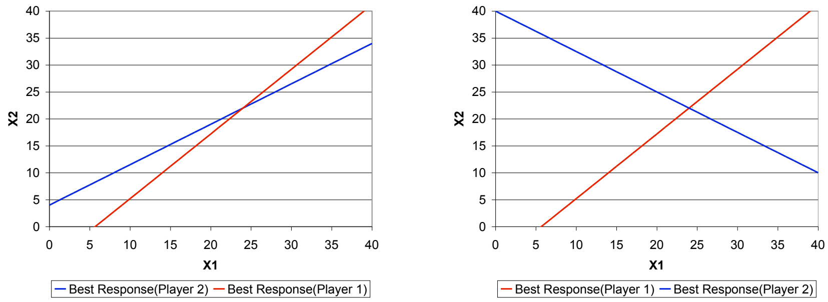

(a) Strategic Complements

(b) Mixed

Figure 1: Best-response functions.

as shown in Figure 2(a). In the Mixed game, however, such a trend is not as clear. If both players select an output that is larger than their respective Nash values in the first round and Cournot best respond to their opponent's choice in the next round, the type-2 player will choose a value that is smaller than the Nash value in the following round. In the following round if they again use the Cournot model, both will choose values lower than the Nash. If this trend continues, in the following round the type-1 player selects a value below the Nash and the type-2 player selects a value above the Nash. If each continues to Cournot best respond, the values slowly converges to the Nash in a way similar to the Strategic Complements. The path in the Mixed game, however, is more chaotic as demonstrated in Figure 2 making it more difficult for subjects in the Mixed treatments to follow.

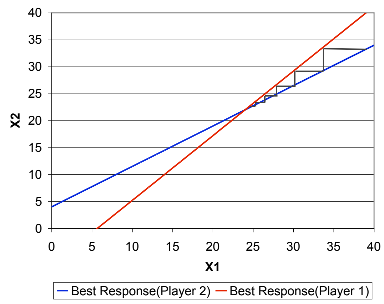

(a) Strategic Complements

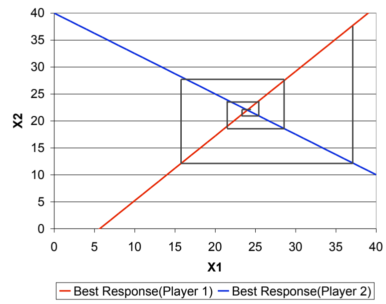

(b) Mixed

Figure 2: Best-response functions showing how the game may evolve using Cournot best-responses.

2.2 *Information*

In addition, there are two different levels of information that a subject receives depending on the treatment. In the first, Low Information, he does not receive any information about what prices other players of his type are playing or the profits these players earn. In each round, he is informed of the value that he chose, the value selected by the player with whom he is matched, and his payoff. In the second treatment, High Information, a player of each type is randomly selected at the beginning of each round to be the public player. At the end of the round, each player sees everything that she would in the Low Information treatment plus all the same information for the public player of her type. Since in every round there is one public player of each type in the High Information treatment, one player of each type gets information about himself each round.

2.3 *Experimental Procedures*

In the experiment, there are 10 players per session who are randomly split into two groups — five type-1 players and five type-2 players. Subjects remain in the same group throughout the experiment. At the beginning of each session, players are assigned to a PC terminal where they are given printed instructions. These instructions include the subject's profit function as well as the payoff function of the other player type. The instructions are then read aloud and subjects are encouraged to ask questions. The instruction period typically lasts around 15 minutes. The instructions for two treatments, Mixed High Information and Strategic Complements Low Information, are included in the appendix in addition to the general computer instructions given to each player before the experiment begins.

At the beginning of each round, each type-1 player is randomly matched with a type-2 player. This helps to minimize the repeated game effects. In each round, subjects are asked to choose an integer price between 0 and 40 and once all subjects have chosen a value, the subject's price, profit, cumulative profit, and the price chosen by the player she was matched with are displayed on the screen. In the High Information treatments, the player also gets all of this information about the public player of his type with the exception of cumulative profit. Throughout the session, players have access to this set of information from all previous rounds including that of the public player in each round in the High Information treatments. After the new prices and profits are displayed, this process is repeated until all 75 rounds have been completed. There are no practice rounds as I am interested in how players learn.

In short, subjects know both their own and their opponent's profit functions as well as having a history of their prices and payoffs from all previous rounds. It is therefore a complete information game where subjects can solve for both best response functions and hence the Nash equilibrium. But I do not know how subjects will use the information they are given or their beliefs about the rationality of others.

There were 16 independent computerized sessions held at Williams College. Four sessions of each treatment were held: Strategic Complements Low Information, Strategic Complements High Information, Mixed Low Information, and Mixed High Information as demonstrated in Table 1.[7] I programmed the experiments in z-Tree (Fischbacher 1999). The subjects were students at Williams College and no one participated in multiple sessions meaning there were 160 subjects. Each session lasted approximately one hour.[8] More detailed data about subject payoffs is available from the author upon request.

[7]There were also three sessions of Mixed Low Information which I did not use. Initially type-2 players had a payoff of $\pi_2 = 1600x_2 + 60x_1 - 20x_2^2 + 13x_1^2 - 30x_1x_2$. I conducted three sessions of the Mixed Low Information game using this payoff equation instead of the that given in Equation 5. This posed problems because type-2 players, even when best responding, were often earning negative profits. As a result, I suspect that some became frustrated during the course of the game and stopped trying to maximize their profits.

[8]The exchange rate was $1 for each 10,000 points. Average earnings for a subject in one of these sixteen sessions was $20. In the first sessions, there was a minimum of $3 that each subject earned regardless of performance. However, this minimum was later increased to $8. In the three sessions of the Mixed game with Low Information where a different payoff function was used, the average payoff was slightly above $11.

3 Hypotheses

After outlining the experimental design, I need to state my hypotheses. They come both from economic intuition and theory, past studies, and computer simulations Professor Gazzale and I ran this summer. To begin, I construct an illustrative model of learning based on the the fEWA learning model outlined in Camerer (2003). This model is a one parameter variation of the experience-weighted attraction (EWA) model developed in Camerer and Ho (1999). In fEWA, strategy probabilities assumed to be are updated after every round and are determined by a weighted average based on how they would have performed in previous rounds if the subject had used that strategy. This model therefore assumes that a player is able to estimate the effectiveness in previous rounds of strategies that she did not select. A player is also assumed to place greater weight on payoffs he actually receives, thus discounting the weight with which they update all others by their "imagination weight." This information and the level at which players discount old information each round varies by subject.

While the fEWA model is ideal for analyzing learning in games, I had to make some modifications to the model in order for it to incorporate the information that subjects received in the High Information treatment. I changed the model so that when a subject sees the public player in a round play value x, he updates the attractiveness of playing x in two ways. First, he calculates what he would have received had he played x, given the value chosen by his opponent and then he discounts this by his imagination weight.

Secondly, he updates the attractiveness of x based upon the profit of the pulbic player in that round. This value carries a weight of one minus the imagination weight.

I used this modified fEWA model and calibrated parameters generated from Chen and Gazzale (2004) to run computer simulations that predicted how subjects would behave. The primary goal of these simulations was to see if the results they generated were logical given the experiment. The simulations predict that, in general, subjects in the Strategic Complements treatments will move to strategies near the Nash equilibrium more quickly than subjects in the Mixed treatments. In addition, they predict that the additional information provided in the High Information treatments will lead subjects to play near the Nash sooner than in the Low Information treatments.

In order state my hypotheses, it is necessary to redefine the definition of Nash equilibrium convergence to better suit a laboratory setting. In theory, a game reaches convergence only if every subject is playing her respective Nash equilibrium value. If a player's choice is even slightly different than the Nash, convergence does not occur. In a laboratory setting with human subjects, this definition is not realistic. I therefore follow the definition used by Chen and Gazzale (2004) and redefine convergence for each of these measures. While the Nash always occurs at (24,22), I call any choice within one of the true Nash value for a player type an ϵ-Nash play as choosing a value within one of the true Nash demonstrates near convergence. For the games in this paper, an ϵ-Nash play for a type-1 player is a value between 23 and 25 inclusive and an ϵ-Nash value for a type-2 player is any choice between 21 and 23 inclusive.

Definition 1 *The level of convergence at round t, L(t), is measured by the proportion of ϵ-Nash equilibrium play in that round. The level of convergence for a block of rounds, $L_b(t_1, t_2) = \sum_{t=t_1}^{t_2} L(t)/(t_2 - t_1 + 1)$, where $0 \leq t_1 \leq t_2 \leq T$ and T is the total number of rounds, t_1 is the first round in the block, and t_2 is the last round in the block. In other words, the level of convergence for a block of rounds is equal to the average level of convergence over that block.*

With this definition, the level of convergence for any round or block of rounds represents the probability that a subject will select an ϵ-Nash strategy. The block convergence measure helps to smooth out variation in convergence across rounds.

In addition to looking at the Nash equilibrium, it is also important to measure how effectively a subject is best responding since it is possible for one player to be best responding with a strategy different than his Nash value. There are many ways to define a variable that measures how effectively a subject best responds to her match in a given round. While it is possible to use a dummy variable that is equal to 1 if the subject is within a given value of the best response (similar to the definition of ϵ-Nash given above), this definition is not entirely satisfying as, after a certain distance, it gives all choices a value of zero even though some may be much closer to a best response than others.

Definition 2 *A player's Best Response Ratio is equal to the her profit divided by the profit she would have earned had she best responded. If she earned a negative profit in a round, her Best Response Ratio for that round is 0.*

I therefore get a variable that has a value varying from 0 when a player's choice is not near a best response and he earns either no profit or a negative profit to 1 when he best responds. This variable provides a more satisfying measure for how effectively a subject responds to his opponent's strategy than any dummy variable.

From the theories presented earlier in the paper, I now form hypotheses about the level of convergence and the Best Response Ratio in the various treatments. I form my hypotheses based on previous experiments, my observations about the experiment and its treatments, and data from the computer simulations of the experiment that I ran.

Hypothesis 1 *For both the High and Low Information treatments, changing from the Mixed treatment to the Strategic Complements treatment significantly increases the level of convergence.*

Hypothesis 1 is based on theories about the characteristics, and best response functions in particular, of the profit functions. If each player best responds to the price his match played in the last round, it is easy to envision the Strategic Complements game moving relatively quickly to the ϵ-Nash equilibrium because of its nice properties whereas the Mixed game will likely take longer since it does not have these same characteristics. These hypotheses are consistent with the laboratory findings in Chen and Gazzale (2004).

Hypothesis 2 *For both the Strategic Complements and Mixed treatments, changing from Low Information to High Information significantly increases the level of convergence.*

Hypotheses 2 is based on the data generated by the computer simulations. It also is consistent with the idea that subjects perform better when they are given more information about the strategies of others as it allows them to develop a better intuition for their payoff function and what subjects of the other type are playing.

Hypothesis 3 *For both the Strategic Complements and Mixed forms, changing from Low Information to High Information significantly increases the subject's ability to best respond.*

Hypothesis 3 comes from the idea that players in the High Information treatments will be able to use the extra information acquired to respond more effectively to opponent strategies than subjects who do not have this additional information.

4 Experimental Results

Figures 3 and 4 display the mean value chosen by type-1 and type-2 subjects in each treatment as well as the Nash equilibrium values for that player type. In addition, the figures present error bars for each mean which represent one standard deviation above and below the mean in that round. The figures show that for both player types in the Strategic Complements treatment, play in later rounds generally remains above the Nash equilibrium. In the High Information treatment, the mean converges to a higher value than in the Low information treatment. For each treatment of Mixed, however, mean play approaches the Nash equilibrium. Additionally, with the exception of

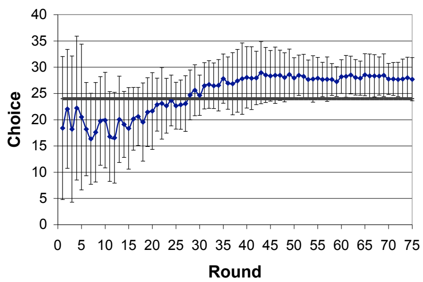

(a) Type-1 Strategic Complements Low Information

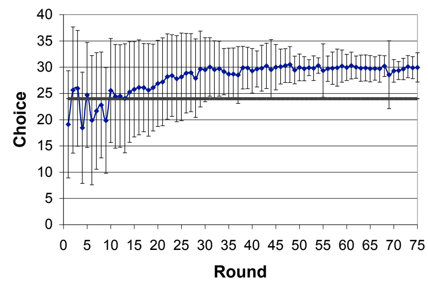

(b) Type-1 Strategic Complements High Information

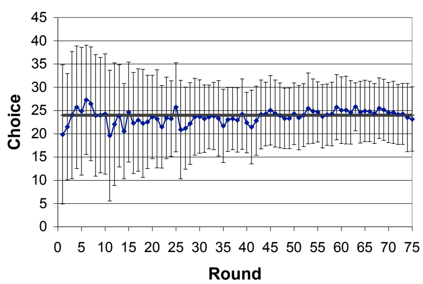

(c) Type-1 Mixed Low Information

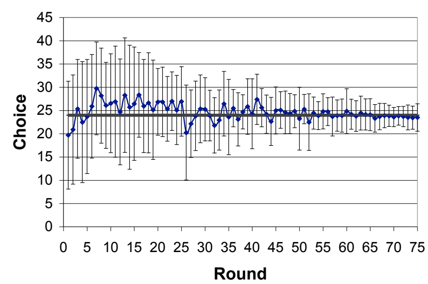

(d) Type-1 High Information

Figure 3: The mean values chosen in each form of the game by type-1 players. The bars represent one standard deviation above and below the mean value in each round. The horizontal line is equal to 24, the Nash equilibrium value for type-1 players.

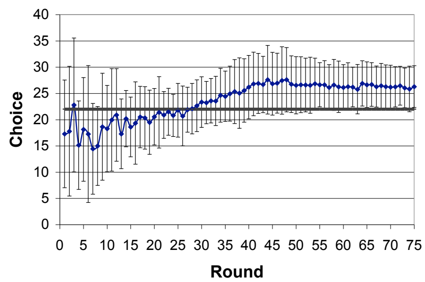

(a) Type-2 Strategic Complements Low Information

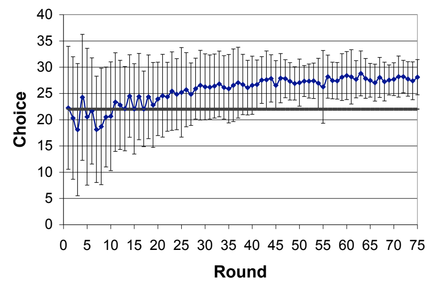

(b) Type-2 Strategic Complements High Information

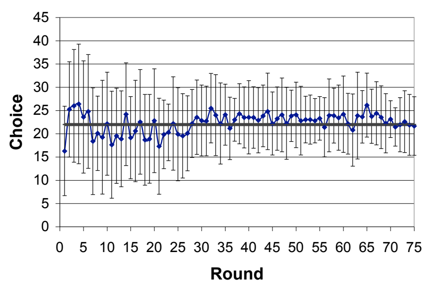

(c) Type-2 Mixed Low Information

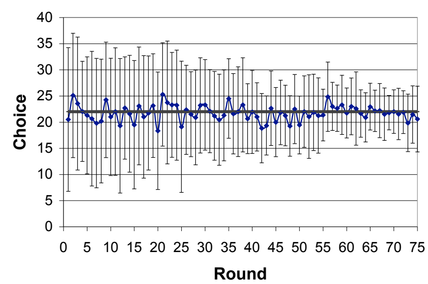

(d) Type-2 Mixed High Information

Figure 4: The mean values chosen in each form of the game by type-2 players. The bars represent one standard deviation above and below the mean value in each round. The horizontal line is equal to 22, the Nash equilibrium value for type-2 players.

the High Information treatment for type-1 players, the figures suggest that the variance is greater in the Mixed game than Strategic Complements for both player types in each treatment.

While these graphs show some general trends, a more rigorous analysis of the data is required. I initially focus on the percentage of players who selected an ϵ-Nash equilibrium value from rounds 51 through 70. This is a good block of rounds to analyze because subjects have had sufficient time to learn about the payoffs and gain a better understanding of the game. By round 50, play is not nearly as erratic as it was earlier in the game. In addition, by stopping at round 70, I hope to avoid any end-game effects that might lead a player to change her behavior.

I look at the probability of ϵ-Nash play for type-1 players, type-2 players, and the randomly matched pairs where both players have selected values within one of their respective Nash. There is a strong correlation between these three values for any given session. This is logical because if each player that if type-1 subjects were consistently playing values around the Nash, type-2 players would be more likely to choose values near the Nash equilibrium as well. The percentage of ϵ-Nash play for each treatment are given in Table 2. Additionally, the percentages for each individual session are given in Table 3 and demonstrate the variance in these values within treatments.

In order to test whether the differences between the probability of ϵ-Nash play in these treatments were statistically significant, I ran permutation tests comparing the mean values for each session and grouped treatments

Treatment	Type-1	Type-2	Both Types
Strategic Complements Low Information	26.0	30.5	10.8
Strategic Complements High Information	3.8	2.3	.5
Mixed Low Information	27.0	21.0	4.0
Mixed High Information	52.5	40.8	23.5

Table 2: Percentage of ϵ-Nash play by treatment for each player type in rounds 51-70. Both Types refers to the percentage that both the type-1 player and her type-2 match for a given round play an ϵ-Nash strategy.

together.[9] Table 4 gives the p-values derived from comparing the ϵ-Nash values of sessions of different treatments.

Result 1 (Level of ϵ-Nash Convergence in Rounds 51-70 for Strategic Complements Versus Mixed)*:*

1. *For type-1 players with Low Information, changing from the Strategic Complements game to the Mixed game does not significantly*[10] *affect the level of convergence.*
2. *For type-2 players with Low Information, changing from the Strategic Complements game to the Mixed game does not significantly affect the level of convergence.*

[9] The permutation test, also known as the Fisher randomization test, is a nonparametric version of a difference of two means t-test. It pools all independent observations and finds a p-value which represents the exact probability of observing a difference between the two treatments like that observed when the pooled observations are randomly divided into two equal-sized groups. The permutation test uses all the information in the sample and therefore has 100 percent power efficiency. For more on the permutation test, see Chen and Gazzale (2004).

[10] When analyzing the data, a significance level of 5 percent or less is referred to as *significant*, while a significance level between five and ten percent is referred to as *weakly significant*. A significance level greater than 10 percent is considered not significant.

Treatment	Session Number	Type-1	Type-2	Both Types
Strategic Complements Low Information	1	11	18	2
Strategic Complements Low Information	2	34	67	22
Strategic Complements Low Information	3	3	1	0
Strategic Complements Low Information	4	56	36	19
Strategic Complements High Information	1	2	0	0
Strategic Complements High Information	2	13	7	2
Strategic Complements High Information	3	0	1	0
Strategic Complements High Information	4	0	1	0
Mixed Low Information	1	19	27	4
Mixed Low Information	2	43	25	10
Mixed Low Information	3	27	18	1
Mixed Low Information	4	19	14	1
Mixed High Information	1	25	16	3
Mixed High Information	2	70	54	39
Mixed High Information	3	53	34	19
Mixed High Information	4	62	59	33

Table 3: Percentage of ϵ-Nash play by session for each player type in rounds 51-70. Both Types refers to the probability that both the type-1 player and her type-2 match for a given round play an ϵ-Nash strategy.

Hypothesis Tests	Player type	p-value
Mixed Low Information < Mixed High Information	Type-1	.0286
Mixed Low Information < Mixed High Information	Type-2	.0571
Mixed Low Information < Mixed High Information	Both	.0286
SC High Information < SC Low Information	Type-1	.0571
SC High Information < SC Low Information	Type-2	.0571
SC High Information < SC Low Information	Both	.1286
SC Low Information < Mixed Low Information	Type-1	.4571
Mixed Low Information < SC Low Information	Type-2	.3000
Mixed Low Information < SC No Low Information	Both	.1714
SC High Information < Mixed High Information	Type-1	.0143
SC High Information < Mixed High Information	Type-2	.0143
SC High Information < Mixed High Information	Both	.0143

Table 4: Hypothesis Tests on ϵ-Nash play in rounds 51-70. "Mixed Low Information < Mixed High Information" means I am testing the hypothesis that the Mixed Low Information treatment has a lower percentage of ϵ-Nash play than the Mixed High Information treatment. The p-value represents the significance level at which we can accept the hypothesis test for the specified player type.

3. *When looking at both player types in the Low Information treatment, changing from the Strategic Complements game to the Mixed game does not significantly affect the level of convergence.*

4. *For type-1 players with High Information, changing from the Strategic Complements game to the Mixed game significantly increases the level of convergence.*

5. *For type-2 players with High Information, changing from the Strategic Complements game to the Mixed game significantly increases the level of convergence.*

6. *When looking at both player types in the High Information treatment, changing from the Strategic Complements game to the Mixed game significantly increases the level of convergence.*

SUPPORT: Table 4 reports the permutation test results for Result 1.

By the findings in Result 1, I reject Hypothesis 1 for all player types in both the Low Information and High Information treatments. These results are unexpected as they demonstrate that the Mixed game is more likely to converge to ϵ-Nash play than the Strategic Complements game in the High Information treatment. This result is not consistent with the results in Chen and Gazzale (2004) which found that supermodularity helped a game converge to ϵ-Nash play.

As demonstrated in Figures 3 and 4, it appears that players of both types in the Strategic Complements treatments often keep playing values above their ϵ-Nash values. In determining potential explanations for this trend, one must look at the different characteristics of the payoffs in each game. In the Strategic Complements game, it is possible for both players to earn a profit above the Nash equilibrium profit whereas this is not possible in the Mixed game. In the Strategic Complements game, total profits are maximized if the type-1 player selects 39 and the type-2 player chooses 36 as their combined payoff will be 8505, or 1325 greater than their total profit if they had played at the Nash equilibrium. In addition, the type-1 player will earn 3978 and the type-2 player will get a payoff of 4527 meaning both players will receive higher payoffs than at the Nash. While in the Mixed game it is also possible for the subjects to earn a higher combined payoff than at the Nash equilibrium, they cannot both earn higher individual payoffs than they would at the Nash. The payoff structure in the Strategic Complements game presents the possibility that players will reach a "collusive" outcome. This term may be misleading as it suggests that players of each type intentionally do not best respond therefore earning profits higher than those at the Nash. Given that subjects were randomly rematched after each round, such intentional collusion would be difficult to sustain. It is possible, however, that subjects of both types were satisfied with the high profits they were earning in these "collusive" situations, and therefore continued satisficing[11] instead of trying to best respond.

Because of the structure of the Strategic Complements game, it is difficult to test for collusion. One possible way to test for such behavior is to check if, in a given round, both the subject and his match earned payoffs higher than

[11]Satisficing is a term coined by economist Herbert Simon to describe a situation where subjects attempt to attain at least a minimum level of a variable (in this case profit) instead of seeking to maximize it.

they would have gotten had they both played their respective Nash values. This is a logical test since it determines whether each player does better in this outcome than if the Nash equilibrium value was played by both subjects. It may be misleading, however, because it is possible that one or both of the subjects involved in the collusion was unaware that they were colluding and instead was simply trying to select the most profitable strategy. This test pooled the data from rounds 51-70 and the results are given in Table 5. In addition, Figure 5 provides the probability of collusive play by round for both the Low Information and High Information treatments in the Strategic Complements game. The data for both forms of the Mixed game are omitted because it is not possible for both players to earn a payoff higher than that at the Nash equilibrium meaning no collusion can occur using this definition.

session type	session number	probability of collusion
Strategic Complements Low Information	1	.60
Strategic Complements Low Information	2	.43
Strategic Complements Low Information	3	.52
Strategic Complements Low Information	4	.34
Strategic Complements High Information	1	.45
Strategic Complements High Information	2	.47
Strategic Complements High Information	3	.73
Strategic Complements High Information	4	.78

Table 5: Probability of a collusive outcome by session in rounds 51-70 of the Strategic Complements game.

As the results from Table 5 demonstrate, players were often earning prof-

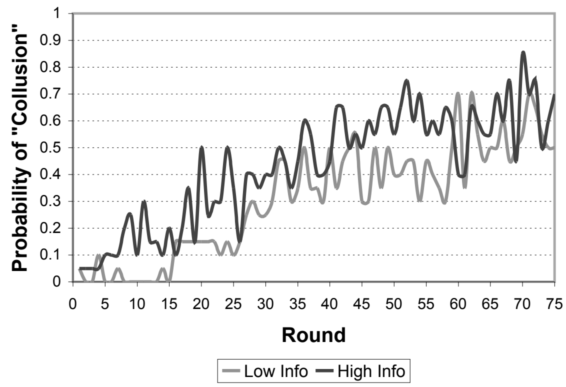

Figure 5: The probability that a pair reaches a collusive outcome in the Strategic Complements game by round.

its greater than those experienced at the Nash equilibrium. In addition, the data suggest that players in the High Information treatment of the Strategic Complements game were more likely to "collude" than those in the Low Information treatment. A possible explanation for this result is that a subject in the High Information treatment sees other players of her type are also playing above the Nash and concludes that it is a best response to her opponent's strategy. This outcome could also occur if a subject is not thinking on the margin. He sees other players of his type best responding and only earning marginally better profits and deems that the costs of the extra effort required to change strategies is greater than benefits of a slightly higher payoff.

Result 2 (Level of ϵ-Nash Convergence in Rounds 51-70 for Low Information Versus High Information)*:*

1. *For type-1 players in the Mixed game, changing from Low Information to High Information significantly increases the level of convergence.*
2. *For type-2 players in the Mixed game, changing from Low Information to High Information weakly increases the level of convergence.*
3. *For players of both types in the Mixed game, changing from Low Information to High Information significantly increases the level of convergence.*
4. *For type-1 players in the Strategic Complements game, changing from Low Information to High Information weakly decreases the level of convergence.*
5. *For type-2 players in the Strategic Complements game, changing from Low Information to High Information weakly decreases the level of convergence.*
6. *For players of both types in the Strategic Complements game, changing from Low Information to High Information does not significantly affect the level of convergence.*

SUPPORT: Table 4 reports the permutation test results for Result 2.

By parts (1), (2), and (3) of Result 2, I accept Hypothesis 2 for all player types in the Mixed game. This result supports the notion that, in the Mixed game, players of both types use the additional information given in the High Information treatment to build a better understanding of the game. This helps them best respond thus increasing the likelihood that the game will move towards the Nash equilibrium and players will select an ϵ-Nash strategy in rounds 51-70.

By parts (4), (5), and (6) in Result 2, I reject Hypothesis 2 for all player types in the Strategic Complements game. This result is unexpected as economic theory predicts that subjects of each type would use the additional information to best respond thus moving the game towards the Nash equilibrium and increasing the likelihood of ϵ-Nash play in rounds 51-70. It is consistent, however, with the possibility that increased information makes subjects more likely to play the "collusive" outcome outlined earlier in the section.

In addition to looking at the probability of ϵ-Nash play in rounds 51-70, I also look at the Best Response Ratios of each treatment over this same period to see if they are consistent with Hypothesis 3. The average Best Response Ratios for each player type by treatment is presented in Table 6. More detailed statistics, which includes information from each session, are provided in Table 7. I ran permutation tests on the data from Table 7 to test Hypothesis 3. The results of these tests are presented in Table 8.

Treatment	Type-1	Type-2
Strategic Complements Low Information	.953	.931
Strategic Complements High Information	.920	.908
Mixed Low Information	.770	.724
Mixed High Information	.859	.845

Table 6: Average Best Response Ratio in rounds 51-70 by treatment and player type.

Result 3 (Best Response Ratio in Rounds 51-70 for Low Information Versus High Information)*:*

Treatment	Session	Type-1	Type-2
Strategic Complements Low Information	1	.972	.977
Strategic Complements Low Information	2	.978	.984
Strategic Complements Low Information	3	.907	.829
Strategic Complements Low Information	4	.956	.935
Strategic Complements High Information	1	.905	.870
Strategic Complements High Information	2	.858	.866
Strategic Complements High Information	3	.949	.932
Strategic Complements High Information	4	.969	.967
Mixed Low Information	1	.827	.841
Mixed Low Information	2	.907	.817
Mixed Low Information	3	.798	.733
Mixed Low Information	4	.549	.504
Mixed High Information	1	.732	.691
Mixed High Information	2	.871	.840
Mixed High Information	3	.919	.921
Mixed High Information	4	.922	.929

Table 7: Average Best Response Ratio in rounds 51-70 by session and player type.

Hypothesis Tests	Player type	p-value
Mixed Low Information < Mixed High Information	Type-1	.1857
Mixed Low Information < Mixed High Information	Type-2	.1286
SC High Information < SC Low Information	Type-1	.1143
SC High Information < SC Low Information	Type-2	.3000

Table 8: Hypothesis Tests on the Best Response Ratio in rounds 51-70. "Mixed Low Information < Mixed High Information" means I am testing the hypothesis that the Best Response Ratio is lower for players in the Mixed Low Information treatment than the Mixed High Information treatment. The p-value represents the significance level at which we can accept the statement in the hypothesis test column for the specified player type.

1. *For type-1 players in the Mixed game, changing from Low Information to High Information does not significantly affect the level of convergence.*
2. *For type-2 players in the Mixed game, changing from Low Information to High Information does not significantly affect the level of convergence.*
3. *For type-1 players in the Strategic Complements game, changing from Low Information to High Information does not significantly affect the level of convergence.*
4. *For type-2 players in the Strategic Complements game, changing from Low Information to High Information does not significantly affect the level of convergence.*

SUPPORT: Table 8 reports the permutation test results for Result 3.

By the findings in Result 3, I reject Hypothesis 3 for each player type in both the Mixed and Strategic Complements games. While it appears unlikely that information improves a subject's Best Response Ratio in the Strategic Complements treatments it is less clear if the information increases a subject's Best Response Ratio in the Mixed game as both p-values are statistically significant at a 20-percent level.

After focusing on Nash play in rounds 51-70, I expand my analysis to focus on all 75 rounds and look at possible explanations for why the additional information in the High Information treatments appears to affect the Strategic Complements and Mixed treatments of the game differently. My initial regressions test how some simple variables such as round, round squared, dummy variables for the type of game and level of information, and interaction variables that multiply each dummy variable by the round affect a subject's Best Response Ratio. The results are shown in Table 9 and some are unexpected. Round and Round^2 behave as I would expect with Round having a positive coefficient and Round^2 having a smaller, negative coefficient. This suggests that as the game progresses, players become better at best responding, but the rate at which they improve declines over the course of the game. The dummy variables provide more interesting results as the regression in Table 9 shows that subjects have more trouble best responding in the Mixed form of the game than in the Strategic Complements form. This is consistent with Chen and Gazzale (2004) in that it suggests that players are better able to best respond to one another and move towards the Nash equilibrium in supermodular games. The dummy variable for High Information, however, is not significant for either type-1 or type-2 players which is consistent with the findings of the permutation test for rounds 51-70 in Table

8. The small coefficient suggests that its lack of statistical significance comes not from a large confidence interval, but instead from its minimal effect on the Best Response Ratio. Additionally, I created a variable that interacted the dummy for High Information with that for the Mixed game so that if the High Information treatments were different but statistically significant for each form of the game, this regression would find the High Information dummy as well as this interaction variable significant. Neither of these was significant, however, suggesting that while the level of information may affect a game's likelihood of reaching the Nash equilibrium, it is not as important in determining a player's ability to best respond.

Additionally, I ran similar probit regressions that used the same variables and focused on the probability that a subject would select an ϵ-Nash strategy. As the results from Table 10 show, many of the variables affect the probability of Nash play similarly to how they influence a subject's ability to best respond. For example, Round is still positive and significant whereas Round^2 and the Mixed dummy variable are both negative. These regressions are also consistent with the results that focus specifically on rounds 51-70 as they find the dummy for High Information to be statistically significant and increase the likelihood of a subject playing an ϵ-Nash value in the Mixed game while decreasing the likelihood of ϵ-Nash play in the Strategic Complements treatments. This is an interesting contrast to the regression results of the Best Response Ratio as information seemingly plays a more important role in effecting ϵ-Nash than a subject's ability to best respond. This statistic suggests that while the additional information may not increase the

Variable	Type-1 Coefficient	Type-2 Coefficient
Round	.01623	.01540
	(.0014)***	(.0014)***
Round^2	-.00013	-.00013
	(.0000)***	(.0000)***
Mixed Game (Dummy)	-.20907	-.17088
	(.0488)***	(.0477)***
High Information (Dummy)	-.01983	-.03311
	(.0462)	(.0504)
Round * Mixed (Interaction)	.00059	.00020
	(.0007)	(.0008)
Round * High Info (Interaction)	.00005	.00069
	(.0007)	(.0008)
Mixed * High Info (Dummy)	.05595	.02415
	(.0530)	(.0525)
Observations	6000	6000
Number of Clusters	80	80
R^2	.2418	.2070

Notes: Robust standard errors in parenthesis are adjusted for clustering.
Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

Table 9: Regression explaining Best Response Ratio. The number below each coefficient represents its standard error. The dummy for Mixed game equals one if the payoff structure of the game is Mixed and zero if it is Strategic Complements. The value of the interaction variables is simply the round times the value of the dummy variable associated with it. The Mixed * High Information dummy variable equals one if the treatment is High Information Mixed and zero otherwise.

Variable	Type-1 Coefficient	Type-2 Coefficient
Round	.00587	.00647
	(.0023)***	(.0022)***
Round^2	-.00008	-.00006
	(.0000)***	(.0000)***
Mixed Game (Dummy)+	-.21168	-.15817
	(.0648)***	(.0608)***
High Information (Dummy)+	-.18745	-.20923
	(.0675)***	(.0584)***
Round * Mixed (Interaction)	.00518	.00244
	(.0015)***	(.0009)***
Round * High Info (Interaction)	.00019	-.00017
	(.0014)	(.0009)
Mixed * High Info (Dummy)+	.27639	.31576
	(.0862)***	(.0967)***
Observations	6000	6000
Number of Clusters	80	80
Pseudo R^2	.0730	.0874

Notes:Coefficients are probability derivatives
Robust standard errors in parenthesis are adjusted for clustering.
+ Signifies coefficient is for discrete change in dummy variable from 0 to 1.
Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

Table 10: Probit results for the probability of an ϵ-Nash play. The coefficients represent the marginal effects of the variable on the probability ϵ-Nash play. The number below each coefficient represents its standard error. The dummy for Mixed game equals one if the payoff structure of the game is Mixed and zero if it is Strategic Complements. The value of the interaction variables is simply the round times the value of the dummy variable associated with it. The Mixed * High Information dummy variable equals one if the treatment is High Information Mixed and zero otherwise.

likelihood of the game moving towards the Nash, it does make a subject more likely to best respond which indicates that players are taking advantage of the information they receive in the High Information treatment.

This additional information may help a subject best respond for a couple of reasons. In the High Information treatment, he sees the choice and payoff of the public player of his type as well as the choice made by this player's match. Such information allows him to gain a better understanding of his payoff function because if the public player earns a high payoff, he knows that she played a profitable strategy against her opponent's choice. Similarly, if the public player earns a low payoff, he knows that her strategy was not strong given her opponent's choice. Hence this additional information gives a subject a better understanding of his payoff function. Additionally, a subject gains valuable information from seeing the choice of another player of the other type. This extra information likely improves his ability to predict what strategy his match will play which allows him to choose a better corresponding strategy.

While these regressions provide some interesting information about how the games evolve, they do not give a satisfying explanation of how players use the additional information provided to make decisions. There are likely other factors which play a key role in determining what choices subjects make. I created a number of other variables and added them to the regressions to see if they helped explain what affects a subject's ability to best respond and play an ϵ-Nash strategy. If the dummy variables used in the initial regressions are not significant in these more comprehensive regressions, it suggests that the new variables are better predictors of the Best Response Ratio and probability of ϵ-Nash play. While such a result seems possible with several of the dummies, it appears less likely with the Mixed dummy variable as there are fundamental differences between the payoff structures that may be difficult to capture in variables. It is also possible that the introduction of new variables will reduce the significance of the Round and Round^2 variables as these effects also ought to be captured in these other variables.

The first of these variables, Instability, measures the standard deviation of the choices made by a subject's opponent in the previous six rounds and the current period.[12] This variable attempts to measure how stable the game has appeared to the subject in recent rounds. In order for a subject to be able to best respond, he must have an idea of what value his opponent will play.

Hypothesis 4 *Instability will be negatively correlated with both (a) the Best Response Ratio and (b) the probability that a subject chooses an ϵ-Nash strategy for both player types.*

Hypothesis 4 is based on the idea that a smaller value indicates that the game has been more stable and predictable over the previous six rounds and the subject is better able to form expectations of what his opponent will play in the current round. With reliable expectations, a subject is more likely to best respond. Similarly, as expectations of opponent play improve, it also seems more likely that the game will move towards the Nash equilibrium thus

[12]Since a subject is randomly rematched with a player of the other type after each round, these choices will likely come from several different opponents.

increasing the probability of ϵ-Nash play.

In the High Information treatment, there are two factors that determine the quality of the information provided by the public player. First, the relative strength of the public player's response is relevant as I would expect a player to benefit more from seeing the public player select a best response than a less profitable strategy. Additionally, the relevance of this additional information is important. If an individual's match in the previous round played a strategy that was greatly different from the public player's match, the additional information is not nearly as valuable as if the two matches played the same strategy.

The Info Quality variable attempts to capture these effects. A subject's Info Quality is comprised of two values multiplied together. Both of these values are derived from the information that she received in the previous round. The first is equal to the Best Response Ratio of her match in the previous round meaning it focuses on how effectively her match in the previous round best responded to his opponent's strategy. This part of the variable is given in the part of Equation 9 labeled Quality. The second value is equal to 40 minus the absolute value of the difference between the choice of the additional player of the other type she saw in the previous round and her opponent's choice in the previous round. This value is given as Relevance in Equation 9.

Player A_1's Info Quality in round n is displayed in the following equation. Assume that in round n, A_1 is matched with B_{1n}. Additionally, in round $n-1$, A_2 was the player of the same type that A_1 saw and A_2 was matched with

$B_{2(n-1)}$ in round $n-1$. Let $q_{n-1}(B_{1n})$ equal the choice of player B_{1n} in round $n-1$.

$$\text{Info Quality}_n(A_1) = \underbrace{BRR_{n-1}(A_2)}_{\text{Quality}} * \underbrace{(40 - |q_{n-1}(B_2) - q_{n-1}(B_1)|)}_{\text{Relevance}}. \qquad (9)$$

Because the data have suggested that subjects use the additional information provided in the High Information treatments depending on if they are playing the Strategic Complements or Mixed game, I created separate variables for each, Info Quality SC and Info Quality Mixed, to include in the regressions. The Mixed form of the variable is equal to Info Quality in Mixed treatments and zero otherwise. Likewise, the Strategic Complements form is equal to Info Quality in the Strategic Complements treatments and zero otherwise.

Hypothesis 5 *The coefficients of Info Quality SC and Info Quality Mixed will be positively correlated with both (a) the Best Response Ratio and (b) the probability of playing a ϵ-Nash strategy for both player types.*

The part of Info Quality labeled Quality places a value on the information that a player gets when she sees the choice and profit of the public player. I expect the Best Response Ratio of the public player in the previous round to be positively correlated with her Best Response Ratio because, if the public player best responds, she will observe his high profits and learn about a profitable under such circumstances. Similarly, if the public player chooses a poor strategy, this is not nearly as helpful as it does not provide her with an example of a successful strategy which she may replicate. The second part of the variable, labeled Relevance, measures how the information relates to the subject's experience. This part is equal to 40 minus the absolute value of difference between the public player's opponent's choice and her match's selection in the last round. I expect a larger value to be more valuable to the subject as it means that her match and the match of the public player chose similar strategies in the previous round. This indicates that the additional information is more relevant and therefore more valuable. Because it appears likely that both parts of this variable to be positively correlated with the Best Response Ratio, the variable ought to be as well. Using the same logic, I also expect the Info Quality to be positively correlated with the probability that a subject plays an ϵ-Nash strategy as this increased information is likely to help players best respond and if both players are best responding, ϵ-Nash play will follow.

The final variable I define is called Surprise and it applies to subjects in both the Low Information and High Information treatments. A subject's Surprise is equal to the square of his opponent's choice in the current round minus the mean of his opponents' choices over the past six rounds, including the current round. Player A_1's Surprise in round n is given in the equation below. Let B_{1n} be A_1's match in round n and let $B_{1(n-t)}$ be A_1's match in round $n-t$ and $q_n(B_{1n})$ equal the choice of player B_{1n} in round n.

$$\text{Surprise}_n(A_1) = \left(q_n(B_{1n}) - \frac{\sum_{t=0}^{6} q_{n-t}(B_{1(n-t)})}{7} \right)^2.$$

Hypothesis 6 *Surprise will (a) be negatively correlated with the Best Response Ratio and (b) will not significantly affect the probability that a subject plays an ϵ-Nash strategy.*

Hypothesis 6 stems from the idea that if a subject's Surprise is large, her opponent's choice is far from what she anticipated. I would therefore expect her to play a strategy that attempts to best respond to the strategy she expected her match to play. Since her match's strategy varies from this expectation, her choice is not a best response. Surprise should not be correlated with the probability of ϵ-Nash play, however, as a player's probability of playing an ϵ-Nash value is not affected by her opponent's decision in the current round because she does not see the value her opponent chooses until after she have already selected a value as well.

Tables 11 and 12 display the results of the regressions which incorporate these new variables. As the data demonstrates, a number of the variables effect the probability of ϵ-Nash play and the Best Response Ratio as the hypotheses predict while there are several that do not.

Result 4 (The Significance of Instability)*:*

1. *For players of both types, increasing the Instability significantly decreases a player's Best Response Ratio.*
2. *For players of both types, increasing the Instability significantly decreases a player's probability of playing an ϵ-Nash strategy.*

SUPPORT: Table 11 reports the regression data for Result 4.

By the findings in Result 4, I accept Hypothesis 4. This result demonstrates the importance of stability in helping a subject best respond to her

Variable	Type-1 Coefficient	Type-2 Coefficient
Round	.00703	.00743
	(.0011)***	(.0015)***
Round^2	-.00005	-.00006
	(.0000)***	(.0000)***
Mixed Game (Dummy)	-.05446	-.03886
	(.0294)*	(.0402)
High Information (Dummy)	-.01962	-.05006
	(.0523)	(.0457)
Round * Mixed (Interaction)	-.00024	-.00080
	(.0005)	(.0007)
Round * High Info (Interaction)	-.00052	-.00016
	(.0005)	(.0007)
Mixed * High Info (Dummy)	.02150	-.10008
	(.0569)	(.0532)*
Instability	-.02299	-.02443
	(.0036)***	(.0029)***
Info Quality SC	.00148	.00175
	(.0011)	(.0008)**
Info Quality Mixed	.00220	.00533
	(.0009)**	(.0009)***
Surprise	-.00131	-.00036
	(.0001)***	(.0001)***
Observations	5920	5920
Number of Clusters	80	80
R^2	.4132	.2983

Notes: Robust standard errors in parenthesis are adjusted for clustering.
Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

Table 11: Regression results for Best Response Ratio. The number below each coefficient represents its standard error. The dummy for Mixed game equals one if the payoff structure of the game is Mixed and zero if it is Strategic Complements. The value of the interaction variables is simply the round times the value of the dummy variable associated with it. The Mixed * High Information dummy variable equals one if the treatment is High Information Mixed and zero otherwise.The other variables are explained in more detail in the text above.

Variable	Type-1 Coefficient	Type-2 Coefficient
Round	.00331	.00258
	(.0026)	(.0022)
Round^2	-.00006	-.00003
	(.0000)**	(.0000)
Mixed Game (Dummy)+	-.13029	-.08792
	(.0650)**	(.0598)*
High Information (Dummy)+	-.06750	-.11816
	(.0619)	(.0540)**
Round * Mixed (Interaction)	.00416	.00192
	(.0015)***	(.0010)**
Round * High Info (Interaction)	.00010	-.00034
	(.0015)	(.0009)
Mixed * High Info (Dummy)+	.02402	.09923
	(.0780)	(.0871)
Instability	-.01282	-.01564
	(.0047)***	(.0041)***
Info Quality SC	-.00384	-.00254
	(.0034)**	(.0013)*
Info Quality Mixed	.00389	.00212
	(.0012)***	(.0011)*
Surprise	.00001	.00001
	(.0001)	(.0001)
Observations	5920	5920
Number of Clusters	80	80
Pseudo R^2	.0863	.1059

Notes:Coefficients are probability derivatives
Robust standard errors in parenthesis are adjusted for clustering.
+ Signifies coefficient is for discrete change in dummy variable from 0 to 1.
Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

Table 12: Probit results for the probability of an ϵ-Nash play. The coefficients represent the marginal effects of the variable on the probability ϵ-Nash play. The number below each coefficient represents its standard error. The dummy for Mixed game equals one if the payoff structure of the game is Mixed and zero if it is Strategic Complements. The value of the interaction variables is simply the round times the value of the dummy variable associated with it. The Mixed * Information dummy variable equals one if the treatment is High Information Mixed and zero otherwise. The other variables are explained in more detail in the text above.

opponent. This result is logical as, in a more stable game, she is more accurately able to form an expectation of her opponent's strategy and select her strategy accordingly. If her opponent's play has varied drastically in recent rounds, she will have trouble best responding even if she knows her best response function because she will not know what value her opponent is going to select. Additionally, it demonstrates that a more predictable game also increases the likelihood that a subject will choose an ϵ-Nash strategy. This result follows from the notion that as stability in a game increases, players of both types are better able to form expectations about their opponents' strategies and best respond accordingly. If players of both types are best responding, the game will move towards the Nash equilibrium and the probability of ϵ-Nash play will increase. While the magnitudes for the coefficients for Instability appear small, it's value varies greatly suggesting that its explanatory power is actually quite large.

Result 5 (The Significance of Info Quality SC and Info Quality Mixed)*:*

1. *Increasing Info Quality SC does not significantly increase a type-1 player's Best Response Ratio.*
2. *Increasing Info Quality Mixed significantly increases a type-1 player's Best Response Ratio.*
3. *Increasing both Info Quality SC and Info Quality Mixed significantly increase a type-2 player's Best Response Ratio.*
4. *For a type-1 player, increasing Info Quality SC significantly decreases his probability of playing an ϵ-Nash strategy.*

5. *For a type-1 player, increasing Info Quality Mixed significantly increases his probability of playing an ϵ-Nash strategy.*

6. *For a type-2 player, increasing Info Quality SC weakly decreases his probability of playing an ϵ-Nash strategy.*

7. *For a type-2 player, increasing Info Quality Mixed weakly increases his probability of playing an ϵ-Nash strategy.*

SUPPORT: Table 11 reports the regression data for Result 5.

Result 5 provides some interesting and unexpected results. Economic theory suggests that both Info Quality SC and Info Quality Mixed should be positively correlated in with both player types in each regression as better information ought to help subjects best respond. The results for the Best Response Ratio regressions appear consistent with this idea as, with the exception of Info Quality SC for type-1 players, both variables have positive, significant coefficients for each player type. While not significant, the coefficient for Info Quality SC for type-1 players is also positive leading me to suspect that with a larger sample size, it would be positive and significant as well. It seems likely that in both the Strategic Complements and Mixed High Information treatments, a subject's ability to best respond depends on the quality of information she saw in the previous round. I therefor tentatively accept Hypothesis 5(a).

In the regression results for the probability of ϵ-Nash play, the two variables behave differently. For both player types Info Quality Mixed remains positive and is at least weakly significant. This is expected as it suggests that a subject who saw a profitable and relevant strategy in the previous round is more likely to play an ϵ-Nash strategy in the current round. Considering that Info Quality is likely to be highest in later rounds when the game is stable and players have a good understanding of their payoff functions, this positive correlation between Info Quality Mixed and the probability of ϵ-Nash play is logical. The coefficient for Info Quality SC, however, is negative and at least weakly significant for both player types. This result implies that a player with a higher Info Quality SC is actually less likely to choose an ϵ-Nash strategy. From Figures 3 and 4 and the data suggesting that players in the Strategic Complements were playing "collusive" strategies, it seems likely that players were getting high Info Quality SC's[13] and then selecting strategies above the ϵ-Nash values. This could lead to Info Quality negatively affecting the probability of ϵ-Nash play. Due to the negative coefficient for Info Quality SC on the probability of ϵ-Nash play for both player types, I reject Hypothesis 5(b).

Result 6 (The Significance of Surprise)*:*

1. *For players of both types, increasing the Surprise significantly decreases a player's Best Response Ratio.*
2. *For players of both types, increasing the Surprise does not significantly affect a player's probability of playing an ϵ-Nash strategy.*

SUPPORT: Table 11 reports the regression data for Result 6.

[13] Even though in these "collusive" outcomes, a player is not best responding to his opponent, he is still earning high profits thus earning a large Best Response Ratio. Additionally, if most players are choosing similar strategies, the public player will have a high Best Response Rate and choose a value close to the other players of his type giving player's a high Info Quality SC in the following round.

By Result 6.1, I accept Hypothesis 6(a). This finding claims that if a subject's match plays a strategy that is different than those played against her in recent rounds, she is less likely to best respond. Such a result is logical as she has an expectation about her opponent's strategy and makes her choice accordingly. If this expectation is wrong, her strategy is less likely to be a best response.

By Result 6.2, I accept Hypothesis 6(b) which states that the value of Surprise does not significantly affect the likelihood that a player will choose an ϵ-Nash strategy. The small coefficients suggest that the lack of significance comes from the fact that the Surprise does not affect the probability of ϵ-Nash play an increased sample size would not change this result. This finding is logical as the value that a subject's match chooses in the current round does not affect the choice of the subject in the same round because he does not see his match's choice until after he has already selected a strategy.

5 Interpretation and Discussion

I expected the introduction of new variables into the regressions would make those used in the earlier regressions statistically not significant as these would provide a more detailed explanation of what caused subjects to best respond and play ϵ-Nash strategies. As the regressions in the previous section show, a number of the original variables become less significant when the new variables were added to the regressions. In some circumstances, however, there were variables which did not lose significance.

In the original regression on Best Response Ratio, Round, Round^2, and the Mixed treatment dummy variable are significant for both player types while all of the other variables are not significant at the 10-percent level. When I add the variables introduced in the previous section, Round and Round^2 remain significant for both player types while the Mixed variable becomes weakly significant for type-1 players and is not significant for type-2 players. Even though Round and Round^2 remain significant in the second regression, the magnitude of their coefficients drops for both player types suggesting that some of their explanatory power is taken by the new variables. In addition, the magnitude of the coefficients for the Mixed dummy variable also decreases in the second regression.

There is also one variable, the dummy for the Mixed High Information treatment, that is not significant in the first regression but becomes weakly significant in the second for type-2 players. This change cannot be attributed to a slight increase in statistical significance as the sign of this variable changes from positive to negative as well. Yet while this result is unexpected, the addition of new variables to the Best Response Ratio regression seems to weaken the explanatory power of the original variables as a whole and provide insight into what specific factors affect a subject's ability to best respond.

These variables have a similar effect on the the probability of ϵ-Nash play regression. In the original regressions, every variable was significant for both player types with the exception of that which interacted Round and the High Information dummy. Many of these variables are no longer significant.

Some of these variables remain significant or weakly significant in the second regressions, such as the Mixed variable dummy and the interaction variable that multiplies Round by the Mixed variable dummy. Yet every variable that was significant in the first regression had a smaller magnitude for both player types in the regression that included the additional variables. This suggests that the new variables play an important role in determining the probability that a subject selects an ϵ-Nash strategy.

Additionally, it is important to note that for both dependent variables, the R^2 and pseudo R^2 values increase by between 18.2 percent and 70.9 percent with the addition of these new variables. It therefore appears that the inclusion of these variables explains how a subject makes strategy decisions more accurately.

The regression results suggest that the most important characteristic in determining what strategy a subject employs, and the effectiveness of this strategy, is stability. The Instability variable clearly captures this idea and, as Tables 11 and 12 demonstrate, this variable is highly significant in determining both a subject's Best Response Ratio and probability of playing an ϵ-Nash strategy. In addition, the magnitude of the coefficient suggest that it plays a major role in determining a player's strategy.

I now look back at the surprising initial results which suggest that High Information makes ϵ-Nash play more likely in rounds 51-70 of the Mixed game while decreasing its likelihood in the Strategic Complements game. If stability is the key component in determining whether subjects choose an ϵ-Nash strategy, I would expect that the mean value of Instability ought to decrease in the Mixed game when switching from the Low Information to High Information treatment and increase in the Strategic Complements game if there is adding this information. As Table 13 demonstrates, the mean Instability values for each treatment and player type in rounds 51-70 are consistent with these results.

	Low Information	High Information
Strategic Complements Type-1	2.19	3.46
Strategic Complements Type-2	2.11	2.30
Mixed Type-1	5.77	4.93
Mixed Type-2	5.39	3.23

Table 13: Mean Value of Instability in rounds 51-70 for treatments and player types in Low Information and High Information treatments.

While I have attempted to explain the primary factors that affect a subject's decision making and contribute to her ability to best respond, there are number of areas which should be explored further. While the second set of regressions made a number of these variables less significant and lowered the magnitude of the coefficient of the others, it did not eliminate the significance of several of the original variables. A more thorough study of the data set and the introduction of new variables or modification of current variables may lessen the impact of these initial variables. In addition, while I looked at the issue of collusion in the Strategic Complements game, I found some evidence suggesting that players were choosing a collusive outcome. This is an area that deserves a closer examination. I have identified Instability as a key variable in determining subject's ability to best respond and allowing the game to progress towards the Nash equilibrium. Yet I have not answered the question of why Mixed game becomes more stable with increased information whereas the Strategic Complements becomes less stable with this additional information. This is a surprising result that also deserves further study.

Appendix 1: Experiment Instructions for Strategic Complements

Low Information Treatment

Experiment Instructions – Treatment SPM-HI-NI

Introduction

- You are about to participate in a decision process in which one of numerous alternatives is selected in each of 75 rounds. This is part of a study intended to provide insight into certain features of decision processes. If you follow the instructions carefully and make good decisions you may earn a considerable amount of money. You will be paid in cash at the end of the experiment.

- *During the experiment, we ask that you please do not talk to each other.* If you have a question, please raise your hand and an experimenter will assist you.

- *After the experiment, we ask that you please do not talk about the specifics of the experiment, especially with those who have not participated in a DM1 experiment.* In addition to helping us maintain experimental control, please know that the specifics of future sessions, including the payoffs to certain choices, will change. You may of course share with others the fact that they can earn cash by entering choices at a computer terminal.

Procedure and Payoffs

- You will be randomly assigned to one of two groups: the Blue group or the Red Group. There will be 5 players in each group. You will stay in the same group for the entire experiment.

- In each of 75 rounds, you will be randomly matched with a player from the other group. You will not know the identity of your Match. Your payoff each round depends only on the decisions made by you and your Match.

- In each of 75 rounds, Red will choose x_1 and Blue will choose x_2. Both of these values will be integers between 0 and 40.

- Red's profit equals $204x_1 + 20x_2 - 18x_1^2 - 15x_2^2 + 30x_1x_2$.

- Blue's profit equals $160x_2 + 60x_1 - 20x_2^2 - 13x_1^2 + 30x_1x_2$.

- Each player's payoff therefore depends not only on what choice they make but also on the choice made by the player they are matched with.

- There will be 75 rounds. There will be no practice rounds. From the first round, you will be paid for each decision you make.

- Your total payoff is the sum of your payoffs in all rounds.

- The exchange rate is \$1 for each 10,000 points.

Information: At the end of each **round**, you are informed of your result of the round:

- Your Choice
- The Choice of your Match for that round
- Your Payoff

We encourage you to earn as much cash as you can. Are there any questions?

Appendix 2: Experiment Instructions for Mixed High Information Treatment

Experiment Instructions – Treatment MIX-HI-RI1

Introduction

- You are about to participate in a decision process in which one of numerous alternatives is selected in each of 75 rounds. This is part of a study intended to provide insight into certain features of decision processes. If you follow the instructions carefully and make good decisions you may earn a considerable amount of money. You will be paid in cash at the end of the experiment.

- *During the experiment, we ask that you please do not talk to each other.* If you have a question, please raise your hand and an experimenter will assist you.

- *After the experiment, we ask that you please do not talk about the specifics of the experiment, especially with those who have not participated in a DM1 experiment.* In addition to helping us maintain experimental control, please know that the specifics of future sessions, including the payoffs to certain choices, will change. You may of course share with others the fact that they can earn cash by entering choices at a computer terminal.

Procedure and Payoffs

- You will be randomly assigned to one of two groups: the Blue group or the Red Group. There will be 5 players in each group. You will stay in the same group for the entire experiment.
- In each of 75 rounds, you will be randomly matched with a player from the other group. You will not know the identity of your Match. Your payoff each round depends only on the decisions made by you and your Match.
- In each of 75 rounds, Red will choose x_1 and Blue will choose x_2. Both of these values will be integers between 0 and 40.
- Red's profit equals $204x_1 + 20x_2 - 18x_1^2 - 15x_2^2 + 30x_1x_2$.
- Blue's profit equals $1280x_2 - 360x_1 - 16x_2^2 + 8x_1^2 - 24x_1x_2$.
- Each player's payoff therefore depends not only on what choice they make but also on the choice made by the player they are matched with.
- There will be 75 rounds. There will be no practice rounds. From the first round, you will be paid for each decision you make.
- Your total payoff is the sum of your payoffs in all rounds.
- The exchange rate is $1 for 10,000 points.

Information: At the end of each **round**, you are informed of your result of the round:

- Your Choice
- The Choice of your Match for that round
- Your Payoff

In each round, the computer will randomly select one Red player and one Blue player. At the end of each round, you are also informed of the result of the randomly selected player of your type. The information presented will be:

- The Choice made by the randomly selected player of your type
- The Choice of the Match of that randomly selected player
- The Payoff earned by the randomly selected player that round

All players of your type will see this same player. It is therefore that possible that the randomly selected player will be you. Players of the other type see a random player of their type that is independent of the random player you see.

We encourage you to earn as much cash as you can. Are there any questions?

Appendix 3: Computer Instructions

Computer Instructions – CM

Announcement Stage

- At the beginning of each round, you enter your Choice.
- You are free to enter any integer between 0 and 40.
- Notice that if you enter a Choice outside of 0 and 40, or do not enter an integer, the computer will tell you that your Choice announcement is not valid and you need to change your selection.

Payoff Calculation

- After all players have submitted a Choice, the computer will calculate your payoff and send this number and other relevant information to your screen.
- This process will be repeated for each round.

Changing Your Entry

- Prior to clicking the **Okay** button, use the **Back Space** key to delete your selection, and then enter your new selection.
- *Once you have submitted your choice, you cannot change it.*

History Box

- At any point in the experiment, you can review all of your previous choices and payoffs by reviewing the **History** box.

- The **History** box will be located towards the bottom of any screen where you are required to make a choice.
- To view rounds that are not visible, use the scroll bar on the right of the **History** box.

Help Box

- At the very bottom of every screen will be a help box entitled **Screen Instructions**.

References

Bohnet, Iris and Richard Zeckhauser, "Social Comparisons in Ultimatum Bargaining," *The Scandinavian Journal of Economics*, 2004, *106*, 495–510.

Bosch-Domenech, Antoni and Nicolaas Vriend, "Imitation of Successful Behaviour in Cournot Markets," *The Economic Journal*, April 2003, *113*, 495–524.

Camerer, Colin, *Behavioral Game Theory: Experiments on Strategic Interactions*, Princeton: Princeton University Press, 2003.

——— **and Teck Hua Ho**, "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, July 1999, *67*, 827–874.

Chen, Yan and Robert Gazzale, "When Does Learning in Games Generate Convergence to Nash Equilibria? The Role of Supermodularity in an Experimental Setting," *The American Economic Review*, December 2004, *5* (94), 1505–1535.

Duffy, John and Nick Feltovich, "Does observation of others affect learning in strategic environments? An experimental study," *The International Journal of Game Theory*, 1999, *28*, 131–152.

Erev, Ido and Amnon Rapoport, "Coordination, "Magic," and Reinforcement Learning in a Market Entry Game," *Games and Economic Behavior*, May 1998, *23*, 146–175.

Fischbacher, Urs, "z-Tree—Zurich Toolbox for Readymade Economic Experiments—Experimenter's Manual," 1999. Working Paper Number 21.

Harrison, Glenn and Kevin McCabe, "Expectations and Fairness in a Simple Bargaining Experiment," *The International Journal of Game Theory*, 1996, *25*, 303–327.

Huck, Steffen, Hans-Theo Normann, and Jorg Oechssler, "Learning in Cournot Oligopoly-An Experiment," *The Economic Journal*, March 1999, *109*, 80–95.

Milgrom, Paul and Chris Shannon, "Monotone Comparative Statics," *Econometrica*, 1994, *62*, 157–180.

——— **and John Roberts**, "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 1991, *58*, 1255–1277.

Potters, Jan and Sigrid Suetens, "Strategic interaction, externalities and cooperation: experimental evidence," February 2005. Working Paper.

Stahl, Dale, "Rule Learning in Symmetric Normal-Form Games: Theory and Evidence," *Games and Economic Behavior*, 2000, *32*, 105–138.