Brain Activity in the Play of Dominant Strategy and Mixed Strategy Games.

By

Rick K. Wilson (Political Science)
Randolph Stevenson (Political Science)
Geoffrey Potts (Psychology)
Rice University

Abstract

We conjecture that the thought processes used to solve dominant strategy games and mixed strategy games are quite distinct. Two-person games with dominant strategies can be treated as simple decision problems that involve no assessment of one's partner. By contrast, two-person games with mixed-strategies require that one think about one's partner. We measure differences in electroencephalogram (EEG) activity while a human subject is playing two-person games. We time-lock the EEG to a common event and use the average across many trials and subjects to find an Event Related Potential (ERP) associated with the common event. The ERP is the brain's response to events -- in this case our different games. Our findings lend support for the idea that subjects respond to types of games differently.

Paper prepared for the Annual Meeting of the Midwest Political Science Association, April 19-22, 2001, Chicago, IL. Enormous thanks are due to Laura Delgado and Denise Torres who assisted in running the experiments reported in this research. Support from the National Science Foundation (SBR-9819943) is gratefully acknowledged. The Foundation bears no responsibility for the content of this research.

"… why does the mathematician not use the culled knowledge of human behavior found in psychology and sociology when formulating his assumptions? The answer is simply that, for the most part, this knowledge is not in a sufficiently precise form to be incorporated as assumptions in a mathematical model. Indeed, one hopes that the unrealistic assumptions and the resulting theory will lead to experiments designed in part to improve the descriptive character of the theory."
(Luce and Raiffa, 1959, p. 5)

Introduction.

Recent work in experimental economics has posed a challenge to traditional views about game theory. While this has not led to a crisis for theorists, the findings have spurred theorists and experimentalists to think hard about what to do to accommodate persistent deviations from normal assumptions about behavior. (Ostrom 1998) has pleaded with theorists to account for cognitive constraints on individuals and to incorporate such limitations into models. Economists like (Camerer 1997)and (Rabin 1998)have argued for "behavioral game theory" in which routine (and consistent) behaviors are incorporated into game theoretic models.

The intent of this research is to take seriously the call for behaviorally informed game theoretic models. We design an experiment to try to peek inside the "black box" of human decision making. We do not offer new game theoretic models. Instead, we propose some fairly simple decision problems and some simple game theoretic tasks and then measure a subject's electrical brain activity while making decisions. Our question is whether new measurement devices available in psychology can help us understand the underlying decision processes of humans in structured settings.

Motivation.

Game theory relies very little on assessing the motives of others. Game theory merely requires that an individual looks internally and decides the best response against what might be played. That is, suppose a rational individual is deciding how to play the game in normal form represented in figure 1. That individual chooses row, while the second player chooses column. In thinking about the game, the row player considers how to play as a column player and then thinks about how to best respond, as a row player, to different column plays. In this game the problem is reasonably simple to solve. The row player has a dominant strategy to play the top row. At the same time, knowing this, the row player understands that the column player will respond by playing right. The joint combination of top row, right column is in equilibrium. In order to solve this particular game, the row player only had to look inwardly and ask what to do in such a circumstance. The other player was hardly necessary in order to determine the appropriate strategy.

<Figure 1About Here>

While game theoretic strategies seem reasonable, there is a great deal of evidence pointing to anomalies in practice. In particular (Holt and Goeree 2001 (forthcoming)) review standard categories of games and conduct a set of experiments designed to test the boundaries of those games. They find consistent violations of standard expectations. It is doubtful that the anomalous behavior is merely a function of the experimental design. Instead it seems to reside in the ways in which subject conceive the problem. These findings by (Holt and Goeree 2001 (forthcoming)) are consistent with a good deal of literature in experimental economics over the past decade. Findings by (Eckel and Grossman 1996) suggest altruistic preferences on the part of subjects in the dictator game. Using a similar game, (Hoffman, McCabe, and Smith 1996) argue that social connectedness may explain some of the pattern of giving that is observed in dictator games. In both instances the puzzling result is why subjects pay any attention to their partner.

More recently a number of models have been offered to explain "other regarding" preferences. Early on, (Rabin 1993) proposed a "fairness" parameter to a utility function enabling him to argue that individuals could account for differences in their partners and respond accordingly. This has given rise to a set of models that account for additional attributes to an individual's utility function that in turn allows for aspects of the partner to be taken into account. (Fehr and Schmidt 1999) point to equity as a primary consideration for an actor, arguing that what a partner gains is a crucial consideration when thinking about allocation problems. The same point is echoed in the model by (Bolton and Ockenfels 2000). (Smith 1998) takes a different tack in his work. Contrary to a model in which individuals account for the relative gain by others, his view on the problem is that individuals are trying to forecast the behavior of others by reading their intentions. He contends that there is an evolutionary advantage to observing another's intention, especially if it leads to mutually beneficial cooperation. There is some support for this position in the sense that subjects in controlled experiments respond to inferences they draw about what their partner intends (see (Burnham, McCabe, and Smith 2000) and (Eckel and Wilson 2001)). Finally, (Sally 2001) offers a model of sympathy that draws on the moral philosophy of Adam Smith (a similar view was expressed by (Frank 1988)). In this model individuals have to put themselves into the shoes of their partners in order to understand their motivation and their strategies.

This speculation by economists and others who draw on game theory has a counterpart in psychology. For psychologists, as well as game theorists, the problem is one of credibly signaling an intention and one of being able to read that intention as credible. Being able to display anger in order to show commitment to a particular action may be useful. But such a display is only valuable if it can easily be read and if it is costly to produce. (Frank 1988) argues that emotions often serve this function for humans in that many emotions are difficult to fake and they are easy to read. Psychologists have been intrigued with one group of people who have the capacity to display emotion, can read emotional states in others, but have a difficult time understanding the content of that emotional state in others. People with autism have a difficult time inferring the intention of other people ((Baron-Cohen 1995) and (O'Connell 1998)).

Baron-Cohen argues that a general “Theory-of-Mind Mechanism” enables people to draw inferences about the intentions of others by going outside themselves. A number of laboratory experiments illustrates the difficulty that subjects with autism and Asperger's syndrome have with reading and interpreting emotional expressions from still photographs. People with autism differ significantly from "normal" functioning individuals who have little difficulty in reading the emotional state of another ((Loveland et al. 1994) and (Baron-Cohen, Wheelwright, and Jolliffe 1997)). This research points to the structure of the brain as crucial for enabling individuals to read signals that give insight into the intention of others. Those without that brain structure find it difficult to read the intention of others and often have great difficulty when negotiating social interaction.

Our sense from this (and other work) is that people pay a great deal of attention to their counterparts in simple games of exchange and negotiation. However, they do so in ways that are different than suggested by game theorists. Under the standard game theoretic model subjects only have to look internally and best respond to what they would play if playing themselves. Our claim, instead, is that people are very aware of their counterpart when choosing a strategy. They try to "read" that counterpart so as to draw an inference about an appropriate strategy and that strategy may not always lead to a Nash equilibrium. Instead, subjects may take a chance that their partner will deviate from best response.

It is impossible to "read" another's mind and determine their underlying motive or intention. However, it is possible to "read" the by-product of thought by picking up electrical activity in the brain. We move to a discussion of this technique and its implications for understanding strategic behavior.

Research Design.

This research uses a design that is unfamiliar to most political scientists. While the task for the subject is fairly standard, the measurement instrument is not. In the experiment a subject makes a series of decisions, either against nature (in the form of rolling a die) or against another person. The matrix of payoffs resembles a game in normal form. Subjects are paid for the decisions they make. While making decisions the subject's brain activity is recorded using a 128 channel device that picks up electroencephalogram (EEG) data from the subject. Subjects make a large number of decisions across five distinct manipulations. Our principle concern is with measuring differences in frontal lobe activity subsequent to the visual stimulus and prior to the key press indicating the subject's choice. A bit different from many studies of Event Related Potentials (ERP), we are less interested in analyzing differences in brains activity following a stimulus and more interested in capturing differences in the decision process prior to a choice. Our "events" then are calculated backward from the action taken by a subject (a response) rather than forward from a stimulus. This section is split into five parts: first it introduces the equipment measuring the subject's EEG; second, it discusses subject characteristics and recruitment; third it details the manipulations; fourth it outlines the procedure used in the experiment; finally it details the hypotheses.

Apparatus.

Brain activity for subjects was acquired with a high density array of electrodes worn by a subject. Electroencephalogram (EEG) data were acquired with a 128 channel Electrical Geodesics system consisting of Geodesic Sensor Net electrodes, Netamps, and Netstation software running on a Macintosh 266 MHz. PPC class computer. EEG data were acquired continuously referenced to the vertex with .1 - 100 Hz. analog filtering and digitized at 250 Hz. The Geodesic Sensor Net is a lightweight elastic thread structure containing plastic pedestals. Each pedestal contains a silver/silver chloride electrode housed in a synthetic sponge. The sponges are soaked in a saline solution to render them conductive. Application of all 128 channels takes approximately 15 minutes.

EEG data were segmented off-line into 1000 msecs. epochs spanning 200 msec. pre-stimulus to 800 msec. post-stimulus. Data were digitally screened for artifact (eye blinks or movements, subject movement, or transient electronic artifact) and contaminated trials were eliminated. Remaining data were sorted by condition and averaged to create the ERPs. Averaged ERP data were digitally filtered at 20 Hz. lowpass to remove residual high-frequency noise, baseline corrected over a 200 msec. prestimulus period, and rereferenced into an average reference frame to remove topographic bias due to choice of reference site (Dien 1998). The subject-averaged ERPs were averaged together to produce the mean waveforms across subjects, the grandaverage waveforms. Statistical analyses were performed on the subject-averaged ERPs with the subject averages being the observations. The waveform and topographic plots and the dipole analyses were performed on the grandaverage data.

All stimulus presentation and behavioral response collection was controlled by an Apple Macintosh 266 MHz. PPC computer running EGIS software (Electrical Geodesics Inc., Eugene). Visual stimuli were presented on an Apple 15" flat-panel active matrix Studio display to reduce 60 - 75 Hz. monitor refresh electrical noise associated with CRT displays. Manual responses were collected with a 4-key microswitch keypad in which on;y the left and right keys were used (Electrical Geodesics, Inc., Eugene, OR).

In all experiments subjects were seated in an adjustable chair with their chin in a chinrest. The chinrest was placed so that subject's eyes were 50 cm. from the center of the flatpanel screen. The chair was adjusted for comfort. Subjects were instructed to remain as still as possible, with their eyes on the fixation mark, throughout the block. Subjects were requested to refrain from blinking as much as possible while the stimuli were appearing. Breaks were provided after every 60 trials (approximately every 5 minutes) so that subjects could rest their eyes.

Subjects.

A total of 13 subjects were recruited for the experiment. Four of the subjects participated as "row" players in the experiment and the remaining nine subjects were "column" players. These latter subjects were fitted with the equipment described above and all recordings were made on these subjects only. All subjects were recruited from the student body at Rice University using an email invitation to a subject pool developed by the first offer. Recordings from subject 7 were anomalous and that subject's readings were excluded from this analysis. Given the remaining eight subjects, five were male and three were female. Subjects were coded for left or right handedness (including a probe for whether their most proximate family members were left handed). All subjects were under the age of 25.

Subjects were informed that they would be paid for their participation and they were told their earnings would depend on the decisions made by themselves and their counterpart. Subjects were told that they could earn as little as $14 (the $10 show-up fee, plus $1 earned in each of the trial blocks) or as much as $90. Subjects were also told that the experiment could last as long as two hours. The longest experiment, from entering the door to leaving lasted one hour and forty-five minutes. Most subjects were done within ninety minutes. Subjects earned between $46 and $72.

Manipulations.

The primary task for subjects involves a visual display followed by a key press in which the subject indicates the choice of a column from a matrix of numbers. Subjects were given five different manipulations and these are displayed on Table 1. The first manipulation was a decision problem with certainty. Subjects were presented the stimulus given in Figure 2. The figure has four small boxes in blue with each box containing a number. The numbers constitute dollar values. There are also four small black boxes with no numbers. Containing these objects are four boxes and each column of boxes are contained in an additional box. At the bottom of the two columns are left and right arrows in red. The subject's choice was between the left or the right column. One additional piece of information was embedded in the stimulus. Between the column boxes, defining a row, were circles. Subjects were told that these circles provided a visual clue to the probability that one row or another would be chosen. In this manipulation a completely filled circle meant that the row would be chosen with certainty. In manipulation A one column always had a larger value then the other column. In addition the numbers in a column were always the same for the top or bottom cell. Moreover, one row was always selected with certainty. In effect this task was a simple decision problem where the subject chose between a smaller and a larger amount of money. As detailed below, subjects were paid according to the choices they made and they knew that their decisions would influence their pay. In this manipulation the column with the largest value was randomized between left and right, as was the row with the probability of being selected with certainty.

<Table 1 About Here>

Manipulation B looked similar to Figure 2 with two exceptions. First, there were different numbers in the top and bottom rows of each column. This meant that a subject could not easily scan between columns and make a simple choice about which column had the largest value. Under this manipulation the two largest values were on a diagonal and the two smallest values were on the opposite diagonal. The largest value was randomized across all cells. The second exception is that several different probabilities were used. Subjects observed a probability for the top row of either .75, .50 or .25 (and the complement of that for the bottom row). Subjects were given the probability for each row by a partially filled-in pie.

<Figure 2 About Here>

Manipulation C presents a visual stimulus similar to that in Manipulation B. Again the two highest values are on a diagonal and the smallest values are on the opposite diagonal. Under this manipulation subjects are told that there is a probability attached to each row. However, they are given two empty circles. This is a decision problem under uncertainty. Subjects are informed that the true probability is being hidden and will only be revealed at the end of the experiment.

Manipulation D resembles Manipulation A in the sense that there is an obvious column choice on the part of the subject. The difference is that a counterpart is added to the experiment and the second subject's payoffs are included on the visual stimulus. The screen for Manipulation D looks similar to that presented in Figure 3. Notice that the counterpart's payoffs are now in the lower right and not shaded. The probability circles are retained, but are left empty to reflect the face that it is uncertain how the counterpart will decide to play -- whether to choose either the top or the bottom row. Manipulation D only involves games with a dominant strategy on the part of the column player. Consequently the column player has an equilibrium strategy. Again, row and column positions were randomly assigned consistent with preserving dominance in the game.

<Figure 3 About Here>

Finally, Manipulation E looks just like Manipulation D. The only difference is that there is no pure strategy equilibrium. Instead all games have a unique mixed strategy equilibrium, albeit one that is difficult for a subject, under time constraints, to calculate.

The aim behind these manipulations is to sort between the presence and absence of a counterpart in the play of a game. Obviously, manipulations A, B and C differ from D and E in that subjects, in the former, play only against themselves or against nature. In the latter manipulations, a counterpart is also making decisions. At the same time there are some comparisons. Manipulations A and D involve very little calculation on the part of the subject -- the dominance relations are obvious. Manipulations B, C and E are similar in that each of them present more complicated decisions have no obvious equilibrium strategy and require thinking about what some other actor might do: either nature or a counterpart. Across all manipulations the design of the visual stimulus and the action by the subject was nearly the same.

Subjects participated in four distinct blocks of decisions. The first block contained 60 distinct decisions that only included Manipulation A problems. The second block contained 120 decisions, 60 of which were from Manipulation A and 60 of which were from Manipulation B. The order of presentation was randomized both within and between the manipulations. The third block contained 120 decisions with 60 drawn from Manipulations B and C. Again the order of presentation was randomized. The final block also had 120 decisions, with 60 from Manipulation D and 60 from Manipulation E. The order of presentation was randomized.

Procedure.

Subjects entered the laboratory and were seated in a chair. They were read a standard introduction and at that point met a second individual who was introduced as the other participant in the experiment. The first subject was told that the other participant would make choices at a computer in an adjacent room. In one block of decisions both individuals would make choices over the same decision tables and they would be jointly paid for an outcome randomly selected by the first participant.

With the preliminaries out of the way the first subject read and signed a consent form. The subject was shown the equipment that would be used and any questions were answered. Two laboratory assistants took cranial measurements, selected a proper sized EEG net, marked the subject's scalp with a grease pencil and then fitted the device. The EEG net was first soaked in a concentrated saline solution and then fitted over the subject's head. The subject was then led to a room equipped with a flat panel monitor and an amplifier that conducted the electrical signals to a computer.

While the electrical readings were being set and the recording program was being started the subject read through a self-paced booklet of instructions concerning how to make decisions. This instruction set is attached as Appendix 1. Once the subject completed the instructions, including answering a set questions probing understanding of the experiment, all last-minute questions were answered. When the subject was ready to begin the experimenter remained in the room and read instructions concerning each block of decisions. Following each block the subject was instructed to draw a poker chip to determine the decision for which the subject would be paid.

Subjects made a total of 420 distinct decisions. As noted above these decisions were made in four blocks. In the first block, a total of 60 decisions were made. In each of the three subsequent blocks subjects made 120 decisions. Subjects were given an opportunity to rest following each group of 60 decisions. For each decision a subject was given 7 seconds in which to make a choice. If a subject failed to make a choice within the allotted time the computer would move to the next decision. If a subject happened to draw a poker chip in which no decision was made another draw was made. Subjects failed to respond in only 27 of the 3220 decisions (with almost half of those failures in Manipulation E and another quarter in Manipulation D). The first player always made a choice between a left and a right column. All of the decisions used the same visual matrix. A subject viewed the stimulus and then responded by pressing a left key or a right key on a separate four-key keypad. While making decisions subjects were cautioned to remain still and to not blink while the game matrix was on the screen.

The second participant in the experiment, a male, was a confederate. Instead of participating, that individual listened to a set of instructions, which the first subject could overhear. Once the first subject was fitted with the equipment and taken to the second room, the confederate was dismissed. Ordinarily a confederate is not used in these kind of experiments. However, in this case we decided to control for all possible stimuli and reasoned that using a confederate was justified because the two subjects could not be in the EEG room at the same time. However, in all decisions a second player's decisions were used. Four subjects were recruited, asked to play as the row player in all games, their choices were recorded, and they were paid at the conclusion of the experiment. During the experiment the participant with the EEG net was randomly paired with one of these four subjects. In this sense the decisions by the first player were always paired with the actions of a second player who was also paid for decisions.

At the conclusion of each block a poker chip was drawn to determine the payment for that block of decisions. Again, the subject was only paid for one decision in the block. Such a form of motivating a subject -- giving large payments for each decision, but paying for only a single, randomly drawn, decision -- works. (Camerer and Hogarth 1999) At the end of the experiment the earnings for the subject were totaled, subjects were shown the table of values to remind them of the payoffs and shown their actions.

Hypotheses.

Our hypotheses are highly speculative and are of two types. The first involves a simple characterization of the cognitive effort that a subject needed to expend. For example, decision problems with certainty (manipulation A) are the least cognitively demanding. All an individual needs to do is compare column numbers. Manipulation B requires a bit more cognitive effort, with subjects now comparing probabilities. This introduces a bit more complexity to the task. While the manipulation resembles a two-person game (this one against nature) it does not require that a person draw any inferences about the behavior of the partner. The second type of hypothesis involves subjects treating games of strategy differently from decision problems. More to the point, we expect that subjects differentiate even among types of games of strategy -- with dominant strategy games treated differently from mixed strategy games.

Hypothesis one expects that there should be a pattern across the response times. If we think in terms of cognitive investment then there is a natural ordering. While decision problems should be the least cognitively demanding, we predict that the two person problem with a dominant strategy should be equivalent to a decision problem with certainty. Therefore we should have: $A < D < B = C < E$

Hypothesis two predicts there will be consistent ERP differences between treatments. We have expectations that ERPs will reflect the simplicity of cognitive task for individuals. This will reflect the response time pattern noted above. In this paper this hypothesis is left unexplored.

Hypothesis three contends there is an ERP difference between decisions under manipulation D and E. This should be captured between 1200 and 800 milliseconds from the key press. Manipulation D decisions involve little cognitive thought and require that the subject pay little attention to the counterpart. By contrast, Manipulation E requires that the subject think hard about the counterpart in the experiment.

Hypothesis four holds that there should be a persistent "oops" wave following manipulation C and manipulation E games. This "Error Related Negativity" (ERN) is predicted at 50 milliseconds following the key press and prior to a distinct cognitive response at around 250 milliseconds subsequent to the key press. (Gehring et al. 1993) Manipulations C and E are similar in that they are both mixed strategy games (although the former is against nature rather then another individual) and so the certainty with which a subject feels s/he has made the correct choice is not clear. However, the magnitude of the "oops" wave is more pronounced in the manipulation E games. As a check, when there are errors in the manipulation A or D games, a very pronounced "oops" wave will appear. This hypothesis is untested in this version of the manuscript.

Analysis.

The analysis is broken in three parts. The first part examines the response times of subjects. The second part looks at overall patterns of electrical activity across treatments. The final part focuses on specific recording regions within very distinct time intervals. Or data present interesting problems. Each subject averages 90 MB of raw data output from the recording device. All data must be filtered and averaged for each treatment. Our readings were reasonably good in that we obtained 37 to 59 clean readings for each 60 decision treatment. With these data comparisons can be offered across treatments by calculating grand means for the pooled subject data.

Response Times

The first cut on the analysis deals with the reaction times for subjects. Part of our conjecture involves an idea that there is a difference in the cognitive capacity necessary to carry out the different decision problems. Response times allow us to get a rough measure of how long subjects took while making decisions and allow us to determine if the task is tackling part of what we intended.

Figure 4 presents box and whiskers plots for each of the treatments. It is quite clear from the figure that the fastest times in making a decision are for the Manipulation A trials. These involve a simple choice between columns, with no other information. Both treatment 1 and 2 involve Manipulation A decisions. They are in different blocks and it is useful to note that subjects take the same response time in both instances. This is so, even though in Block 2 subjects have to recognize the difference between certain and probabilistic decisions. A similar phenomenon occurs when comparing decisions with risk. The response times for probabilistic decisions are similar even though they are randomized with other types of decisions across blocks. All-in-all this gives us confidence that our treatments hold up independently with what they are paired.

Hypothesis 1 held that the relationships across manipulations should be:

$$A < D < B = C < E$$

However, it is clear from "eye-balling" the data that this does not hold up. Our thought was that the simple decision trials (A) would be less than the dominant strategy trials (D), which in turn would be less than the probabilistic (B) and uncertain (C) decisions, which in turn would be less than the mixed strategy (E) trials. Except for the dominant strategy trials (D) the ordering is preserved. However, the response time for the dominant strategy trials is considerably higher than any of the trials that only require simple decision making. What is not clear from these data is whether the difference between the simple decision making problems and the game theoretic problems is due to task complexity or to thinking about the other player's moves. It could be that the problem is more complex because subjects must digest another four numbers. On the other hand understanding dominant strategies in these games is not altogether different from solving the simple decision problem.

<Figure 4 About Here>

Global Patterns

The primary concern in this analysis is to capture aspects of the decision process. At best we can be concerned with general patterns of electrical activity preceding an action (the key press indicating a choice). Our approach to this problem is a bit different from standard practices in ERP studies that look at signals subsequent to the introduction of a stimulus. We are concerned with events prior to an action. However, none of the actions are of equal length (something that Figure 4 makes clear). We look backward from the key press. On average, humans take 800 milliseconds to execute an executive motor control function. As a consequence we focus on events happening just prior to the subject starting a finger press. We assume that the decision is made just prior to starting the finger press. This means looking at events ranging from 800 to 1200 ms. prior to the recorded key press.

Figure 5 reproduces the EEG data collected for all channels, averaged across all subjects, for 1200 ms. prior to the key press and 240 ms. following the key press. The figure compares Manipulations D and E. These data are set to a 200 millisecond baseline that occurs just prior to the recording of the key press. While we are not interested in looking across all channels, this figure provides some sense as to the complexity of the data. However, several points are worth making. First, it is clear that our channels are recording in a manner that is consistent with other ERP findings. The motor control functions exhibit a characteristic pattern that is defined by a slow downturn that goes negative, with a sharp spike upward around 120 ms. following the key press. The motor functions are recorded by the central band of channels. Figure 5 is oriented such that the top of the figure is the front of the brain. Channels 30, 31, 32, 13, 7 on the left and 113, 119, 107, 112, and 106 on the right all exhibit a standard pattern, with the ERP occurring when it should. Figure 6 pulls out and expands four channels, two on the left and two on the right. Each shows a characteristic pattern of a slightly declining, relatively smooth decrease in discharge at around 800 ms, with a fully negative discharge at 50 ms. and then a spiking upward at 200 ms. This result gives us considerable confidence that the EEG recordings were made in a proper manner. We see the same patterns in all of the other treatments.

<Figures 5,6 About Here>

What stands out in these data are the strong, long term, positive charges emerging at channels associated with the prefrontal cortex. In particular these are channels 127, 22, 26, 27 on the left and 126, 125, 14 and 8 on the right. The organizing and decision making functions are located in these areas. At the same time these channels tend to be a bit noisy, reflecting eye movement, etc. Despite this the overall patterns are consistent with those observed from other ERP studies. The question now is whether the patterns appearing in these data point to consistent differences between the two manipulations with human interaction.

Local Patterns

To test whether there are differences across the dominant and mixed strategy games (hypothesis three) it is necessary to take a clustering of channels, a time slice of interest and then calculate averages for the treatments. We identified channel groupings on the scalp above the medial, left temporal, and right temporal prefrontal cortex and examined the waveforms from these channels for likely ERPs and for differences between treatments. Because we were expecting the key differences to occur just prior to activating the motor functions (a finger press), we focused on a time slice that ranged from 1050 ms to 700 ms before the recorded key press.

Examining the wave forms in manipulations D and E, one can immediately identify a potentially large difference between manipulations in many channels of the medial prefrontal cortex. This is especially evident from Figure 7, which illustrates the differences between the dominant strategy and mixed strategy wave forms. Of particular interest are the clear jumps in one wave form or the other in channels 127, 126 and 8. Amplitudes in this area between 1148 ms. and 440 ms before the response appear to be consistently lower in manipulation D than in E.

<Figure 7 About Here>

We also observe a less obvious, but consistent, pair of positive ERPs across the medial prefrontal channel group. This can be seen in channels 17, 18 and 15 on Figure 8. However, the jumps are seemingly only associated with Manipulation E and not

Manipulation D. As well, these ERPs seemingly arise late -- appearing very close to the onset of motor control.

<Figure 8 About Here>

Differences in the left and right temporal prefrontal cortex channel groups appear to be less consistent across channels but are in some cases pronounced, especially in the range from 960 ms. to 540 ms before the response. In addition, the polarities of the ERPs in this range seem to vary temporally.

These impressions were tested with simple difference of means tests, in which average amplitudes for each time-slice for each of our eight subjects provide the observations. Tables 2 through 4 show that there are significant differences across an array of channels between the dominant and mixed strategy games. These differences hold up under a number of different time slices, with almost all of those times reflecting a period in which a choice is being finalized. Moreover, these channels are in the temporal prefrontal cortex, which is consistent with the area that is associated with higher order decision making.

<Table 2-4 About Here>

Table 5 compares the activity between the left and right sides of the brain. Here the signs of the differences indicate a pattern that lends credence to Hypothesis three and the more general point to this research. It appears that there is a bias toward positive readings on the right side of the brain under the dominant strategy games. This makes sense in that most of the calculating structure for the brain resides on that side. Meanwhile, the mixed strategy games, which require attention to one's partner, have a stronger signal on the left side of the brain. Speculation holds that social decisions (weakly) dominate this side of the brain.

Conclusion.

This is a very preliminary study aimed at determining whether an ERP methodology has anything useful to contribute to the study of game theory. We hesitantly think so. Our hesitation stems from the fact that we have not completely digested these data and that we are uncertain whether the experimental design allows us to make the inferences we wish. The former will be resolved with more analysis.

We have learned several things from this pilot study. The first is that the ERP methodology requires very clearly delineated experiments. The experimental design for such an experiment requires uniformity across all manipulations.   Although we expected that our design retained the same visual patterns, we suspect that the addition of the second player's payoffs added considerable complexity to the experiment. This is evident in the differences in response times. We are heartened, however, in that in the fourth block of experiment that we find differences between two manipulations, even though both have the same additional level of complexity. Second, we have learned that the ERP methodology requires very specific hypotheses concerning the regions of brain activation. Moreover, that methodology works much better as an "event" activated from a known stimulus. Third, we suspect that higher order cognitive tasks, which include strategy decisions in light of another player, are going to be difficult to capture with this methodology. It will certainly require careful thought and consideration when designing experiments. Finally, there are severe data problems associated with this methodology. There is a lot of data and it requires a number of tools to get to the point of conducting analysis. At this point we have only scratched the surface.

On the positive side, the ERP methodology is relatively non-invasive. Subjects can be fitted with the equipment relatively quickly and there is no associated discomfort. The equipment is relatively inexpensive and can be used with only modest training. Unlike PET or fMRI, subjects can be run rather quickly. Although we ran only nine subjects in this pilot test, many ERP studies now use several hundred subjects. Also on the positive side is the fact that an enormous amount of data is generated that allows researchers to cut at underlying processes in many different ways.

The promise of this particular study is that we find evidence that leads us to think that subjects are approaching games differently from what game theorists might predict. If true, this gives us tremendous insight into the boundaries on human cognition. We think this holds some promise for better focusing our game theoretic models. We stake out a modest position, however, and we will not be certain until we analyze these data (and others) more completely

Bibliography

Baron-Cohen, Simon. 1995. *Mindblindness: An Essay on Autism and Theory of Mind.* Boston: MIT Press.

Baron-Cohen, Simon, Sally Wheelwright, and Therese Jolliffe. 1997. Is There a 'Language of the Eyes'? Evidence from Normal Adults, and Adults with Autism or Asperger Syndrome. *Visual Cognition* 4 (3):311-331.

Bolton, Gary E., and Axel Ockenfels. 2000. ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review* 90 (1 (March)):166-193.

Burnham, Terence, Kevin McCabe, and Vernon L. Smith. 2000. Friend-or-foe intentionality priming in an extensive form trust game. *Journal of Economic Behavior & Organization* 43 (1):57-73.

Camerer, Colin F. 1997. "Progress in Behavioral Game Theory.". *Journal of Economic Perspectives,* 11 (4):167-88.

Camerer, Colin F., and Robin M. Hogarth. 1999. The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Journal of Risk and Uncertainty* 19 (1-3):7-42.

Dien, Joseph. 1998. Issues in the application of the average reference: Review, critiques, and recommendations. *Behavior Research Methods, Instruments, & Computers* 30 (1):34-43.

Eckel, Catherine C., and Philip Grossman. 1996. Altruism in Anonymous Dictator Games. *Games and Economic Behavior* 16:181-191.

Eckel, Catherine C., and Rick K. Wilson. 2001. Why Fairness?: Facial expressions, evolutionary psychology, and the emergence of fairness in simple bargaining games. Rice University.

Fehr, Ernst, and Klaus M. Schmidt. 1999. A Theory of Fairness, Competition and Cooperation. *Quarterly Journal of Economics* 114 (3):817-868.

Frank, Robert. 1988. *Passions Within Reason: The Strategic Role of the Emotions*. New York: W. W. Norton & Co.

Gehring, W.J., B. Goss, M.G.H. Coles, D.E. Meyer, and E. Donchin. 1993. A neural system for error detection and compensation. *Psychological Science* 4:385-390.

Hoffman, Eliazbeth, Kevin McCabe, and Vernon L. Smith. 1996. Social distance and other-regarding behavior in dictator games. *American Economic Review* (1996):653-60.

Holt, Charles A., and Jacob Goeree. 2001 (forthcoming). Ten Little Treasures of Game Theory, and Ten Intuitive Contradictions. *American Economic Review*.

Loveland, Katherine A., Belgin Tunali-Kotoski, Deborah A Pearson, and Kristin A Brelsford. 1994. Imitation and expression of facial affect in autism. *Development & Psychopathology* 6 (3):433-444.

O'Connell, Sanjida. 1998. *Mindreading: An Investigation into How We Learn to Love and Lie.* New York: Doubleday.

Ostrom, Elinor. 1998. A Behavioral Approach to the Rational Choice Theory of Collective Action. *American Political Science Review* 92 (1):1-22.

Rabin, Matthew. 1993. Incorporating Fairness into Game Theory and Economics. *American Economic Review* 83 (5):1281-1302.

Rabin, Matthew. 1998. Psychology and Economics. *Journal* of Economic Literature (36):11-46.

Sally, David. 2001. On sympathy and games. *Journal of Economic Behavior & Organization* 44 (1):1-30.

Smith, Vernon L. 1998. The Two Faces of Adam Smith. *Southern Economic Journal* 65 (1):1-19.

Table 1
Listing of Manipulations (Number of Trials per Subject in Parentheses)

Decision Problem with Certainty	Decision Problem with Probability	Decision Problem with Uncertainty	Two-person Problem with Dominant Strategy	Two-person Problem with Mixed Strategy
A (120)	**B** (120)	**C** (120)	**D** (60)	**E** (60)

Table 2
Medial Prefrontal Cortex:
Mean amplitude from 1024 ms to 928 ms before response

Channel	#17	#18	#15	#9	#24	#19	#16	#3
avg. difference between D and E	-2.408	-3.171	-3.248	-2.652	-2.386	-2.693	-2.562	-3.197
t-test	-0.988	-1.021	-1.438	-2.070	-1.945	-1.424	-1.284	-1.525

Table 3
Medial Prefrontal Cortex:
Mean amplitude from 1148 ms to 440 ms before response

Channel	#17	#18	#15	#9	#24	#19	#16	#3
avg. difference between D and E	-1.416	-2.163	-2.48	-2.05	-1.481	-1.826	-1.71	-2.557
t-test	-0.874	-0.889	-1.392	-1.660	-1.448	-1.245	-1.126	-1.504

Table 4
Medial Prefrontal Cortex:
Mean amplitude from 1148 ms to 1048 ms before response

Channel	#17	#18	#15	#9	#24	#19	#16	#3
avg. difference between D and E	-1.951	-3.19	-2.818	-2.528	-2.31	-2.651	-2.426	-3.266
t-test	-0.593	-0.794	-0.975	-1.370	-1.577	-1.068	-0.928	-1.193

Table 5
Temporal Prefrontal Cortex:
Mean amplitude from 960 ms to 540 ms before response

Channel	#127	#126	#22	#14	#26	#8	#33	#1
avg. difference between manipulation D and E	-0.95	3.99	-0.59	-1.45	-1.15	3.2	-2.83	5.54
t-test	2.724	3.069	-3.226	-2.388	-3.244	3.956	-1.840	-0.329
temporal location	Left	Right	Left	Right	Left	Right	Left	Right

Figure 1
Simple Two-person Game in Normal Form

	Left	**Right**
Top	(5,2)	(4,5)
Bottom	(3,4)	(3,3)

Figure 2
Sample Screen for Manipulation A

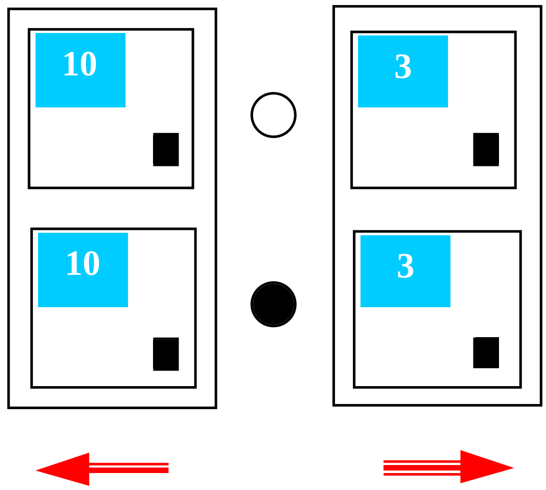

Figure 3
Sample Screen Manipulation D

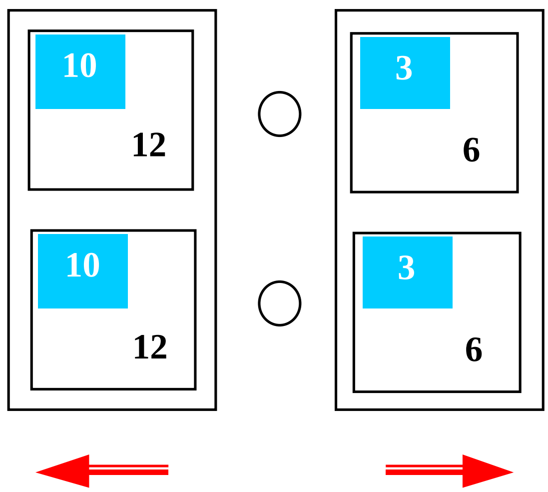

Figure 4
Response Times by Treatment

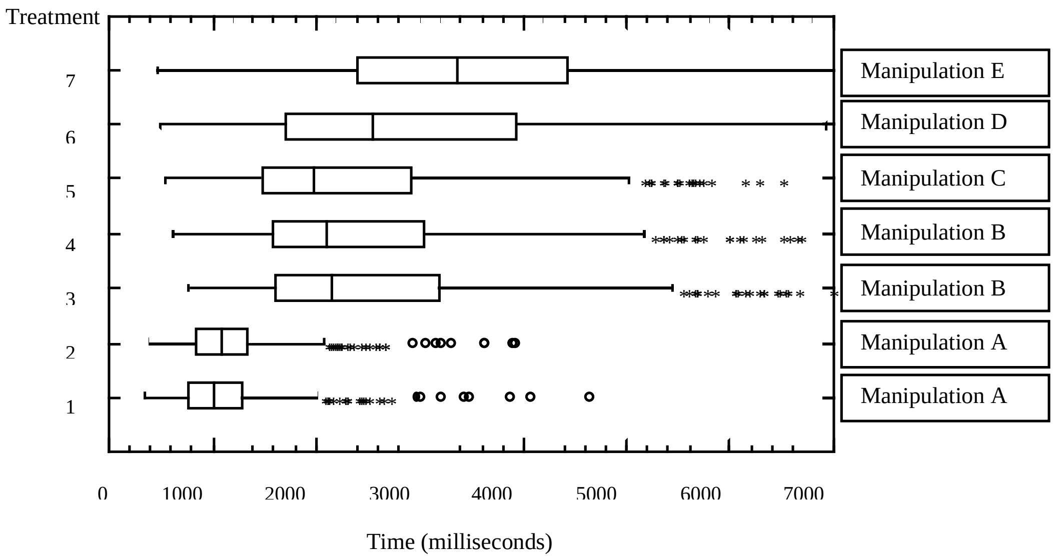

Figure 5
Map of EEG Recordings -- All Channels for Manipulations D and E

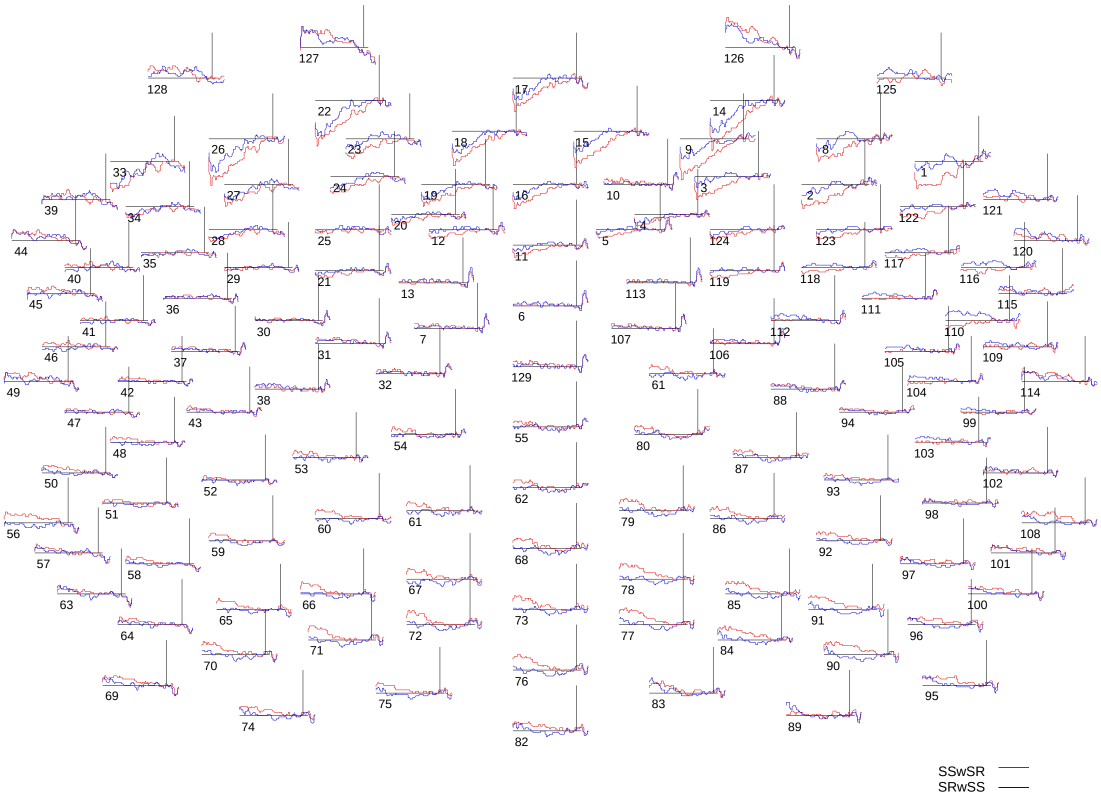

Figure 6
Motor Response -- Selected Channels

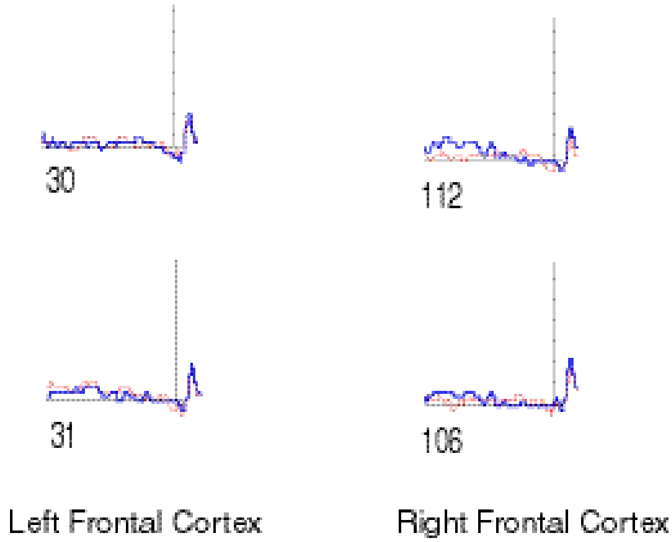

Figure 7
Left and Right Temporal Prefrontal Cortex -- Selected Channels

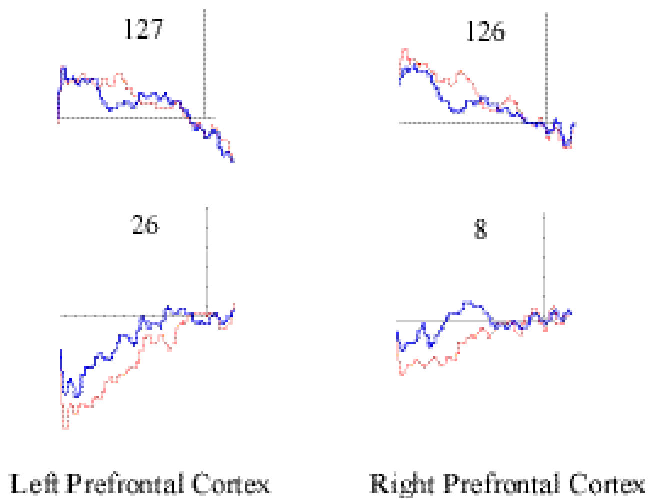

Figure 8
Medial Prefrontal Cortex -- Selected Channels

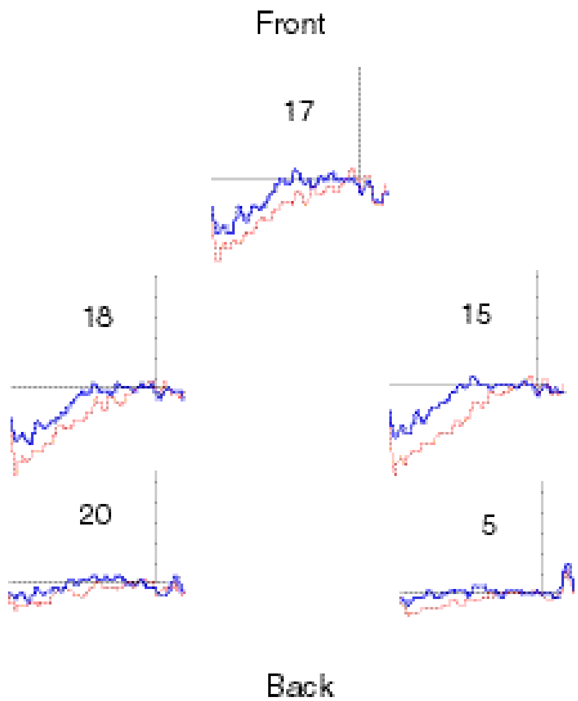

Appendix 1

Screen 1

Welcome to the Experiment

In this experiment you will face a number of decision tasks. There are two types of tasks. In one set you will make a series of individual choices. In the second set you will make choices jointly with another person.

During the course of the experiment you will have a limited time to make your decisions. As well, during the experiment we will be measuring your surface brain activity.

Screen 2

The next several pages will illustrate the types of decisions you will make. All of the pages you will see (and the screens you will see on the computer) will be very similar.

Some of the screens will contain more information then others.

Screen 3

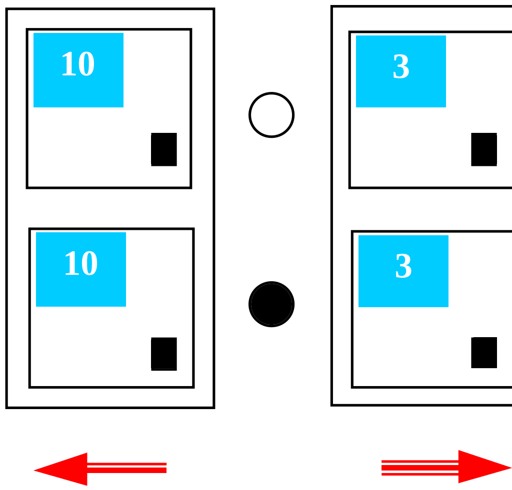

On this screen you see two large boxes making two columns, two medium-sized boxes inside each column representing a top or bottom row, blue-shaded boxes with numbers and black squares. There are also circles and arrows.

The number in each blue-shaded box represents dollars and is what you would earn is a specific column and row is chosen.

In every decision you can choose either the left column or the right column, To make a choice you will press the key marked “1” for the left column or the key marked “4” for the right column.

The circles to the right of the boxes indicate the likelihood the row will be chosen. That likelihood is represented by the degree to which the circle is filled in. In this example the bottom row always selected because its circle is completely filled in (and the corresponding circle for the top row is empty).

Screen 4

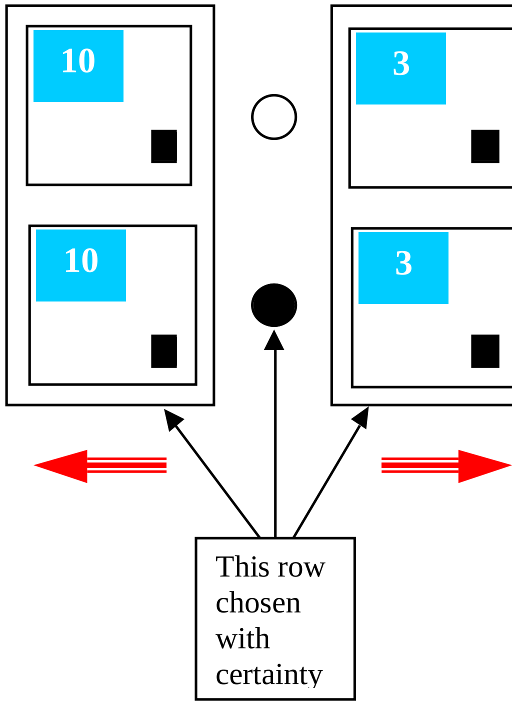

This screen shows you that, because the bottom circle is completely filled in, the bottom row will be chosen. Your choice is simply to pick either the left or right column.

Screen 5

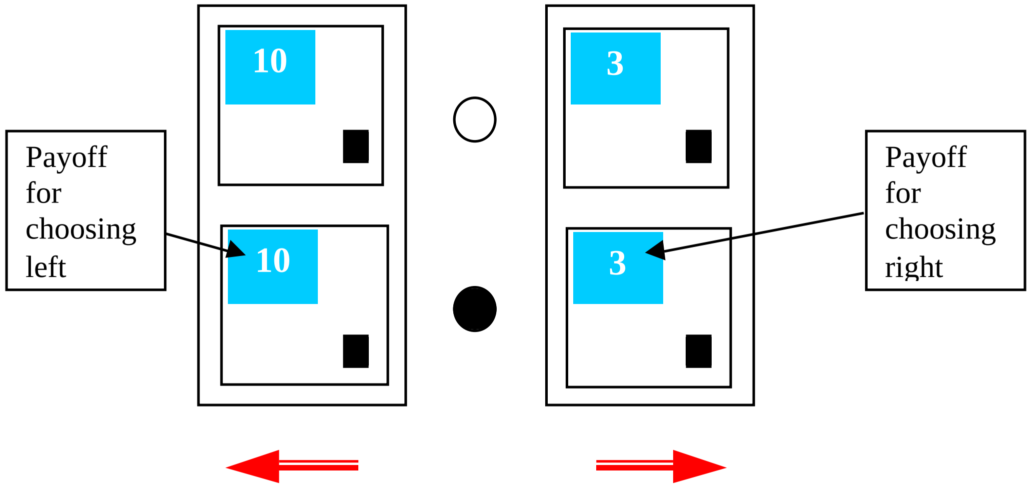

Because the bottom row will be chosen, with certainty, you will earn either the left or right payoff, depending on YOUR choice.

Screen 6

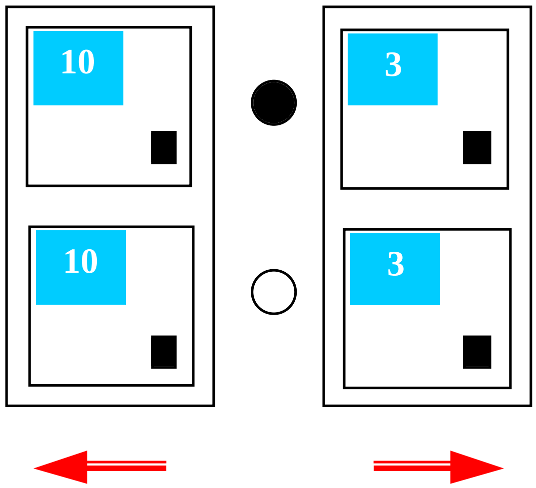

In this example, if you pressed the “4” key, point out to the experimenter how much you would earn.

Screen 7

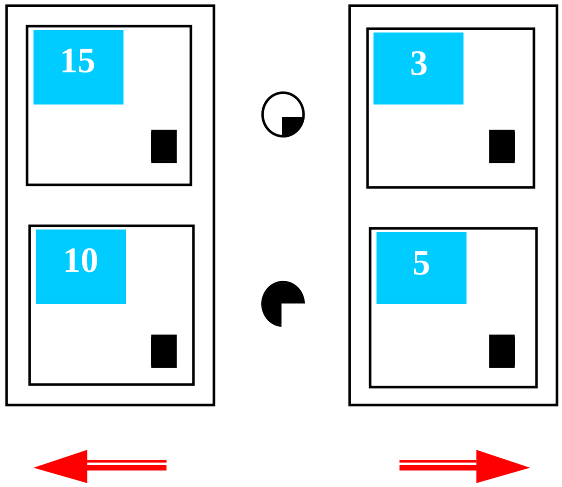

On this screen you see a different type of decision. There are different numbers in each of the boxes. Again your choice is between the left column of boxes and the right column of boxes.

As you can see the likelihood of obtaining the top row is about 1/4 while the likelihood of the bottom row is about 3/4 (you can tell this by how much the two circles are filled). The roll of an 8-sided die will select the top or bottom row. If a 1 or 2 is rolled the top row will be selected. If a 3 through 8 is rolled, the bottom row will be selected.

Screen 8

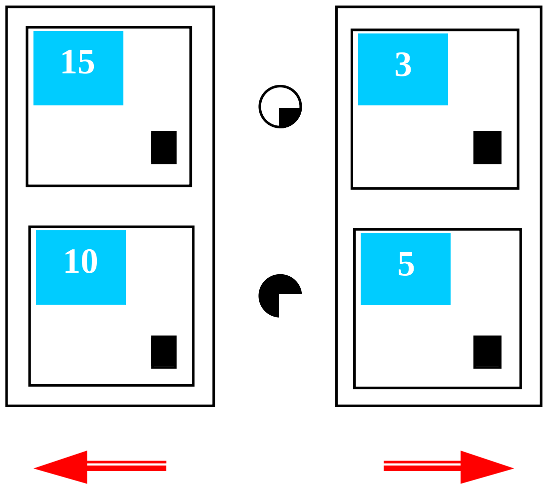

Suppose you pressed the “1” key, rolled the 8-sided die and a 2 turned up. Point out to the experimenter the payoff you would receive.

Screen 9

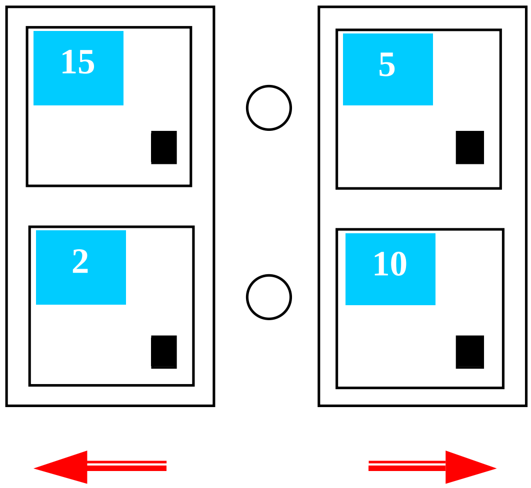

On this screen notice that the numbers are different from the other screens. Again your choice is the left or right column. This screen is also different in that both of the circles are unfilled. This means that the likelihood that one row or the other is selected is unknown. There is a likelihood for each row, but you will not be told what it is.

Screen 10

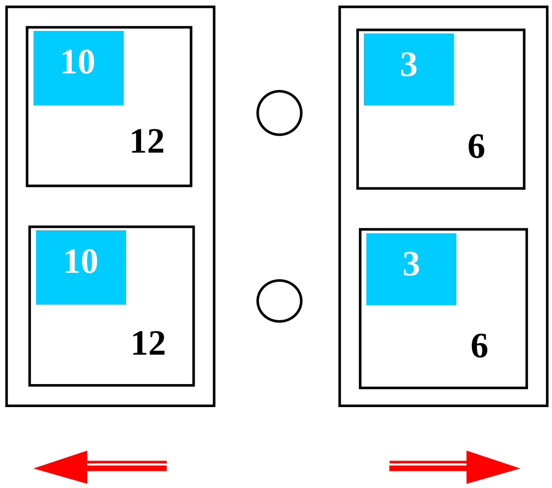

This screen is more complicated because it introduces another decision maker. Your decision is still to choose the left or right column. The other decision maker will choose the top row or the bottom row.

Your earnings are in white and are in the boxes that are shaded in blue. The other decision maker's earnings are given in black and in the lower right of each box.

Screen 11

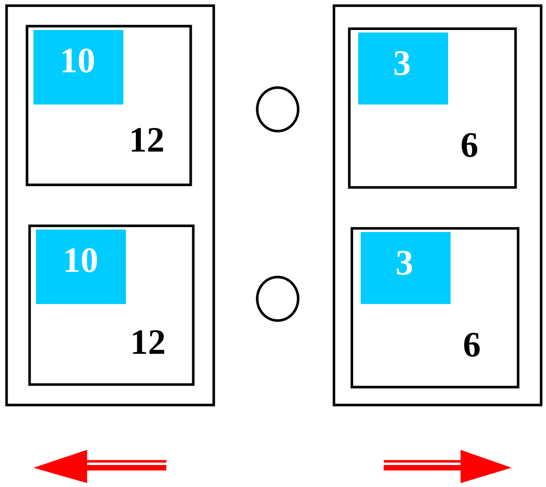

The other decision maker will be making choices at the same time (but you will not be told that choice until later), the circles to the right are not filled.

In this example, suppose you pressed the “1” key and the other decision maker chose the top row.

(1) Please point out to the experimenter your payoff.

(2) Please point out the other decision maker’s payoff.

Screen 12

You have now seen all the types of screens you will see in the experiment. Your potential earnings will always appear in the same place.

You will make a large number of decisions. These will be broken out in 15 minute blocks. You will not be paid for each decision. Instead, at the end of each 15 minute block, you will draw a poker chip to determine the decision for which you will be paid. Each decision is equally likely to be picked, so treat each one very carefully. Once the decision is selected you will be told the choice by the other decision maker or you will be told to roll a die – depending on the type of decision. You will then be paid, in cash, for that decision.

If you have any questions, please ask them now. If you would like, please feel free to re-read any of the instruction sheets.