Political Psychology, Vol. 27, No. 3, 2006

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

Rick K. Wilson
Rice University

Randolph Stevenson
Rice University

Geoffrey Potts
Rice University

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 modest support for the idea that subjects respond to types of games differently.

KEY WORDS: Event-Related Potential, Strategic Behavior, Behavioral Game Theory, Methodology

Recent work in experimental economics has posed a challenge to traditional views about game theory. While this has not led to a crisis, the findings have spurred theorists and experimentalists to think hard about how 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 Rabin (1998) and Camerer (2003) have argued for "behavioral game theory" in which routine (and

Published by Blackwell Publishing. Inc., 350 Main Street, Malden, MA 02148, USA, 9600 Garsington Road, Oxford, OX4 2DQ,
and PO Box 378 Carlton South, 3053 Victoria Australia

consistent) behaviors are incorporated into game theoretic models. In some ways this can be seen as formal theory rediscovering psychology.

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 examine simple game theoretic tasks and measure a subject’s electrical brain activity while making decisions. Our question is whether physiological measurement 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 look internally and decide the best response against what might be played. For example, 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 row player has a dominant strategy to play the top row. 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.

While game theoretic strategies seem reasonable, there is a great deal of evidence pointing to anomalies in practice. In particular Holt and Goeree (2001) review standard categories of games and 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 subjects conceive the problem. These findings by Holt and Goeree (2001) are consistent

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

Figure 1. Simple Two-person Game in Normal Form.

with a good deal of literature in experimental economics over the past decade. For example, Eckel and Grossman (1996) suggest altruistic preferences on the part of subjects who can give some of their own earnings to others in the "dictator" game. In the "trust" game Eckel and Wilson (2005) detail the ways in which the counterpart's attractiveness makes a difference in trusting and trustworthiness. In both cases the puzzling result is why subjects pay any attention to their partner.

A number of models have been offered to explain "other regarding" preferences. Rabin (1993) proposed a "fairness" parameter in a utility function that allows individuals to account for differences in their partner's intention and respond accordingly. This has given rise to models that detail additional attributes of an individual's utility function that allows for considerations of the partner. Fehr and Schmidt (1999) point to equity as a primary consideration, arguing that what a partner gains is crucial when thinking about allocation problems (see also Bolton & Ockenfels, 2000). Smith (1998) takes a different tack, arguing that individuals try 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. 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 partner in order to understand their partner's motivation and strategy.

An imaginative experiment by McCabe, Houser, Ryan, Smith, and Trouard (2001) uses fMRI (functional Magnetic Resonance Imaging) to look at subject brain activity when paired with either another human or a computer. Looking at a sequential game, with the "imaged" subject in both the first and second mover position, they find significant differences in levels of brain activation in the prefrontal cortex when subjects are paired with another individual or paired with a computer. However, these findings only hold true for subjects who attempt a cooperative move in the sequential game. For subjects who can be labeled non-cooperative, there are no such differences. McCabe et al. (2001) speculate that this offers evidence that subjects, who are trying to play non-Nash strategies, pay attention to their counterparts.

Our sense from this 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 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 principal 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 brain 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 (Electrical Geodesics Inc., Eugene, OR) running on a Macintosh 266 MHz. PPC class computer (Apple Computer, Cupertino, CA). 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.

The electroencephalogram (EEG) is the scalp-recorded electrical activity due to excitatory and inhibitory post-synaptic potentials in the dendritic arborizations summed across large populations of neurons in laminar (cortical) tissue. In general higher frequency, lower amplitude EEG is associated with more neural activity or arousal. Most of the scalp-recorded EEG reflects "background" neural activity, not associated with specific cognitive processes. The brain activity associated with specific mental operations is very small compared to the background EEG and must be enhanced, usually by signal averaging. If the same event is repeated multiple times in the same condition and multiple segments of the EEG are time-locked to that event, those segments of EEG can be averaged together. Since the background EEG is random with respect to the event, it converges to zero as the number of segments increases, leaving the brain's electrical response to the event, the Event-Related Potential or ERP (electrical potential reflects the potential for current to flow between two points).

ERPs are typically digitally sampled at between 250 and 1,000 Hz, providing a theoretical temporal resolution of 1–4 ms, although realistically, individual ERP components take tens or hundreds of milliseconds to resolve. Never the less, the temporal resolution of the ERP is in the millisecond range, the temporal scale of mental operations. In contrast, the spatial resolution of ERPs is relatively poor, with a limit somewhere in the centimeter range. While the advent of dense-sensor array recording (roughly 64 electrodes or more) has improved the spatial resolving power of ERPs, locating the neural sources of ERP signals remains ambiguous.

The head is a volume conductor of electrical signals. EEG recording electrodes are located on a surface of that volume. Estimating the number, location, strength, and orientation of sources of a signal from recordings made on the surface of a volume conductor of that signal is what is known as an inverse problem. Inverse problems, as a class (including acoustic and thermal inverse problems, as well as electrical), are ill-posed, i.e., they do not have a unique solution. Thus any source-localization of an ERP signal will have a resolution, at best, in the centimeter range and will have an inherent ambiguity.[1]

In contrast, hemodynamic methods, e.g., positron emission tomography (PET) and function magnetic resonance imaging (fMRI), have a spatial resolution in the millimeter range and are a true volumetric measurement, i.e., there is no inverse problem and thus no ambiguity about the collection of signal sources. However, hemodynamic measures are an indirect index of neural activity; they do not reflect the action of neurons, but rather the delivery of blood-borne metabolic fuel for those neurons. The temporal resolution of hemodynamic methods is limited by the time-course of the hemodynamic response, which is roughly 6–12 seconds, and the speed of the recording technology. In this design we opt for the coarser, but faster, EEG data.

The EEG data used here were segmented off-line into epochs spanning 1,200 ms pre- to 240 ms post-response. 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.

---

[1] For a general discussion of EEG biophysics see Nunez and Srinvasan (2005); for ERP methods see Handy (2004); and for reviews of ERPs in cognitive research see Rugg and Coles (1997) and Zani and Proverbio (2002).

lowpass to remove residual high-frequency noise, baseline corrected over the first 200 ms of the epoch, and rereferenced into an average reference frame to remove topographic bias due to choice of reference site. 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.

In all experiments subjects were seated in an adjustable chair with their chin in a chinrest. The chinrest was placed so that subjects' eyes were 50 cm. from the center of the flat-panel 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 five minutes) so that subjects could rest.

Subjects

A total of 13 subjects were recruited for the experiment. Four of the subjects participated as "row" players in the experiment, made their choices in private, and were never introduced to the other subjects. The remaining nine subjects were "column" players, and they were fitted with the equipment described above. All subjects were recruited from the student body at Rice University using an email invitation to a subject pool developed by the first author. Recordings from subject 7 were anomalous and 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 45 minutes. Most subjects were done within 90 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

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

	No Counterpart	Strategically Irrelevant Counterpart	Strategically Relevant Counterpart
Certain Outcome	**A** (120)	**D** (60)	
Probabilistic Outcome	**B** (120)		
Uncertain Outcome	**C** (60)		**E** (60)

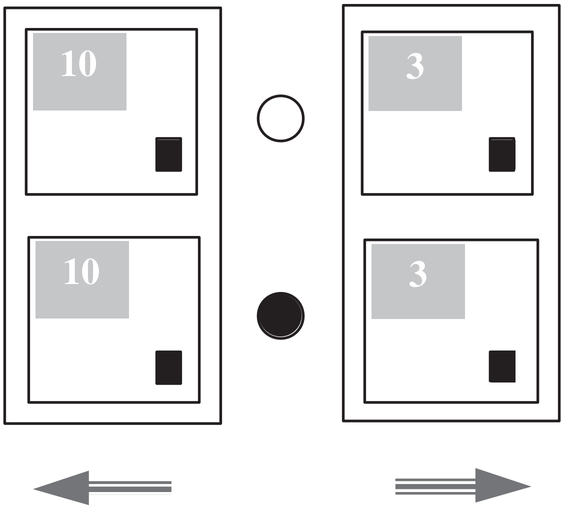

Figure 2. Sample Screen for Manipulation A.

blue with each box displaying a dollar value. There are also four small black boxes with no numbers. Surrounding these objects are four boxes, and each column of boxes is 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 right column. One additional piece of information was embedded in the stimulus. Between the column boxes, defining a row, are circles. Subjects were told that these circles indicated the probability that one row or another would be chosen. A completely filled circle meant that the row would be chosen with certainty. In manipulation A one column always had a larger value than the other column, the numbers in a column were always the same for the top or bottom cell, and one row was always selected with certainty. This task was a simple decision problem where the subject chose between a smaller and a larger amount of money. 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.

Manipulation B was 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

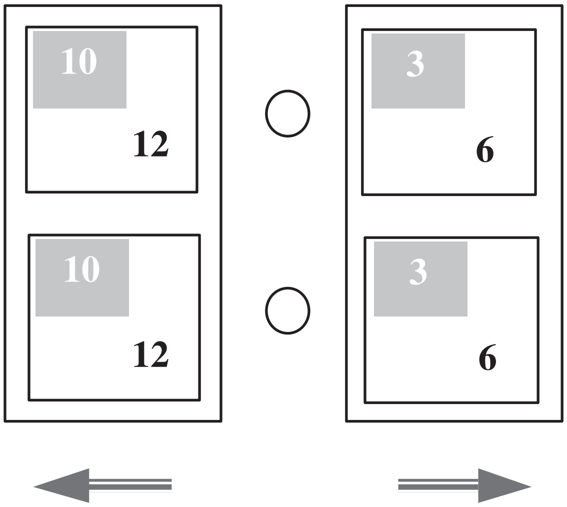

Figure 3. Sample Screen Manipulation D.

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.

Manipulation C presented 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 hidden and will only be revealed at the end of the experiment.

Manipulation D resembles Manipulation A in 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 fact 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.

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. Moreover, subjects know that they are playing against another human.

Our primary concern with 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 involved in making decisions. These manipulations also allow us to make other comparisons. It may be that what we observe is simply a function of task complexity. Manipulations A and D involve almost no calculation by the subject—the dominance relations are obvious—so these tasks should look similar if there is no effect due to the presence of a counterpart. As well, manipulations B, C, and E involve complicated decisions with equilibrium that are difficult to calculate (and impossible to do with manipulation C because no information is revealed concerning Nature's move). The cognitive demands are much higher, irrespective of the presence of a human counterpart.

Subjects participated in four 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. Again, 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 (a confederate) 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 available from the first author on request. Once the subject completed the instructions, including answering 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 moved 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 not to 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. The decisions that involved a counterpart were in fact made by a second player. Prior to conducting the experiment, 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 (see Cameron, 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

We have two measures of different aspects of the decision process: the response times, and the electrical activity of the left and right prefrontal cortex. If we find that both measures match for two treatments, then we cannot reject the hypothesis that the cognitive processes underlying decision making in the two treatments are the same. Here we assume that the level of uncertainty associated with a decision problem (certain, probabilistic, and uncertain) captures the complexity of the task and would not, of course, expect response times and prefrontal activity to be the same for decisions of different complexity. We make explicit "null" hypotheses for no-difference between treatments with and without a counterpart, controlling for task complexity.

> H1: Response times will be the same for simple decision problems (A) and simple strategic problems (in which the counterpart's actions are irrelevant for a player's best response—D).
>
> H2: Response times will be the same for simple decision problems with uncertain outcomes (C) and outcomes that are uncertain because they depend on an uncertain choice of the other player (E).
>
> H3: Activity in the prefrontal cortex for decisions involving a counterpart (D and E) will be the same as those that do not involve a counterpart (A, B, and C), controlling for task complexity.

Most of the discussion takes place with respect to a period of cognition extending from 800 to 500 milliseconds prior to the press of the key. This is a period in which individuals complete their choice and prior to beginning motor control functions.

Analysis

The analysis is broken into 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.

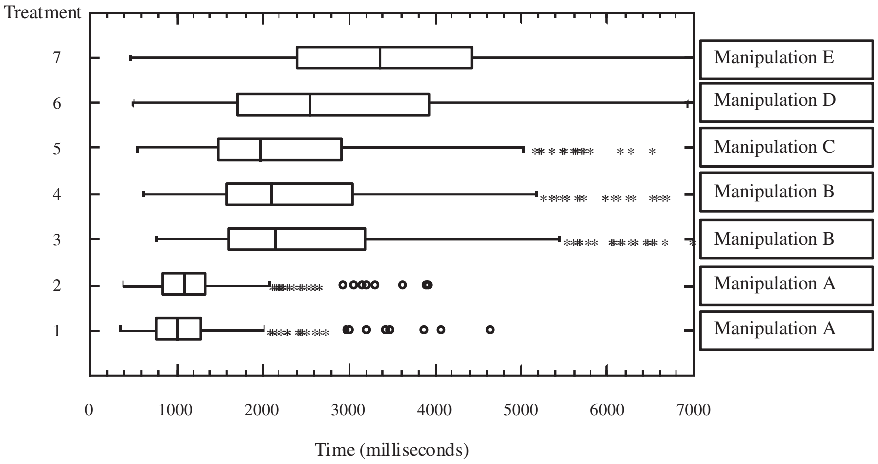

Figure 4. Response Times by Treatment.

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. Even though they are in different blocks, it is useful to note that subjects take the same response time in both instances ($t = 1.76$, $df = 918$, $p = .08$). 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 manipulation B decisions are similar even though they are randomized with other types of decisions across blocks ($t = 1.05$, $df = 918$, $p = .30$). All in all this gives us confidence that our treatments hold up independently despite being paired with other types of decisions.

Hypothesis 1 holds that equivalent types of decisions ought to reflect the same response time. However, it is clear from "eye-balling" the data that this does not hold up ($t = 22.13$, $df = 918$, $p < .001$). We reject the hypothesis that response times are the same for manipulations A and D. Despite the strategically equivalent nature of these decision problems, subjects take longer to determine, in the presence of a counterpart, that they should simply choose the higher numbers. This could be due to their thinking about the counterpart, but may of course result from differences in the complexity of the task that are not accounted for by the level of uncertainty. Consequently, we cannot take this for unequivocal evidence that the counterpart is the source of the difference.

We also reject the second hypothesis that response times are the same for treatments C and E ($t = 11.90$, $df = 918$, $p < .001$). Again, these are strategically similar choices, because for both the subject is uncertain whether the row player will choose up or down. In the unlikely event that the subject could calculate the mixed strategy of her opponent, then the strategic structure of E mirrors B rather than C. However, the response times for B and C are quite similar, and both are different from E. In a simple ANOVA on response times across these three treatments (and controlling for order effects) there is a difference ($F_{(2)} = 88.98$, $p < .001$). However, a post hoc test of manipulations B and C indicates that we cannot reject the null hypotheses that the response times are the same ($F_{(1,1376)} = 2.30$, $p = .13$).

Global Patterns

The primary concern in this analysis is to capture aspects of the decision process in the presence or absence of a human counterpart. Our approach 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 500–600 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 500 to 800 ms. prior to the recorded key press.

Figure 5 isolates EEG data collected for four selected channels, averaged across all subjects. The time interval extends from 1,200 ms prior to and 240 ms after the key press. We use these channels to benchmark our findings with those common in other ERP studies. These channels cover the primary motor cortex anterior to the central sulcus. Motor control recordings exhibit a characteristic pattern in which the signal goes slowly negative followed by a sharp spike upwards around 150 ms following the key press. On the figures the vertical line indicates the key press. Each channel shows a characteristic pattern of a slightly declining, relatively smooth, decrease in discharge at around 600 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.

If we focus on channels specific to the prefrontal cortex we find patterns similar to those reported in other ERP studies. These channels are selected because

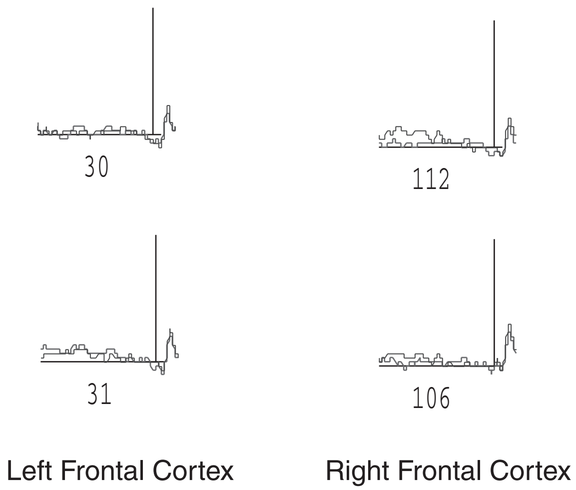

Figure 5. Motor Response—Selected Channels.

executive brain function related to decision making has long been associated with the prefrontal cortex (see Pribram & Luria, 1973; Stuss & Knight, 2002). The critical question 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 between simple decision problems and strategic problems, we focus on clusters of channels over the prefrontal cortex and examined the waveforms from these channels for differences between treatments. Because we expect the key differences to occur prior to activating the motor functions (a key press), we focused on a time slice that ranged from 800 ms to 500 ms before the recorded key press. Averages for each subject first are calculated for all decisions within each treatment and block. Those averages are then normalized (with a mean of 0 and standard deviation of 1) for each subject across all treatments. This is to take into account individual differences (which were considerable). Normalized scores, then, are used in all subsequent analysis and all subjects are pooled. We take a long enough time slice so as to capture the final decision prior to motor control functions taking over. Once we have the data structured in a useable manner, we then find the average value across the two prefrontal cortex channel clusters noted above (keeping track of the right and left hemisphere). Averages and standard deviations are included in Table 2.

Hypothesis 3 suggests that there should be no difference between Manipulation A and D. Using a t-test we can reject that hypothesis (t = 6.30, df = 1,798,

Table 2. Descriptive Statistics of Prefrontal Cortex Channel Clusters Broken Out by Manipulation and Hemisphere—500–800 ms. from Keypress (All Data Have Been Normalized to Each Subject)

Left Hemisphere					
Manipulation					
	A	**B**	**C**	**D**	**E**
Average	.011	−.317	−.289	−.265	.488
Standard Deviation	.868	.782	.564	.894	.739
N	1,200	1,200	600	600	600
Right Hemisphere					
Manipulation					
	A	**B**	**C**	**D**	**E**
Average	.237	−.318	−.156	−.267	.454
Standard Deviation	.904	.633	.702	.656	.555
N	1,200	1,200	600	600	600

$p < .001$ for the left, prefrontal cortex and $t = 12.13$, $df = 1{,}798$, $p < .001$ for the right, prefrontal cortex). However, as Figure 6 shows, the differences are not enormous. Figure 6 is a box and whiskers plot of the distribution of averaged values over the period 500–800 milliseconds prior to the key press. All five manipulations are plotted on the figure, with the left and right hemispheres plotted separately. In this figure Manipulation A, with the simplest decision task, has higher positive values than decisions under Manipulation D which is the simpler of the two games of strategic choice. The fact that Manipulation D is more negative implies pre-frontal cortex activity in an area implicated in social control. However, as can be seen from the figure, the differences between Manipulation A and D are smaller than some of the others.

Hypothesis 3 also suggests that there should be no difference between treatments B, C, and E. A simple ANOVA allows us to reject that hypothesis ($F = 273.39$, $df = 2$, $p < .001$ for the left hemisphere and $F = 303.48$, $df = 2$, $p < .001$ for the right hemisphere). However, using a Bonferroni test over all pairwise comparisons, we do not find that all three manipulations significantly differ from one another. Manipulations B and C are quite similar. There is no difference between them for the left hemisphere ($t = –.78$, $df = 1{,}798$, $p = .43$) but there is a difference in the right hemisphere ($t = –4.95$, $df = 1{,}798$, $p < .001$). By comparison both the left and right hemisphere values for the prefrontal cortex channels are significantly different between manipulations B and E and C and E. The t-statistics range between −16 and −25 for the left and right hemispheres. These differences can be seen from the box and whiskers plots on Figure 6. Here

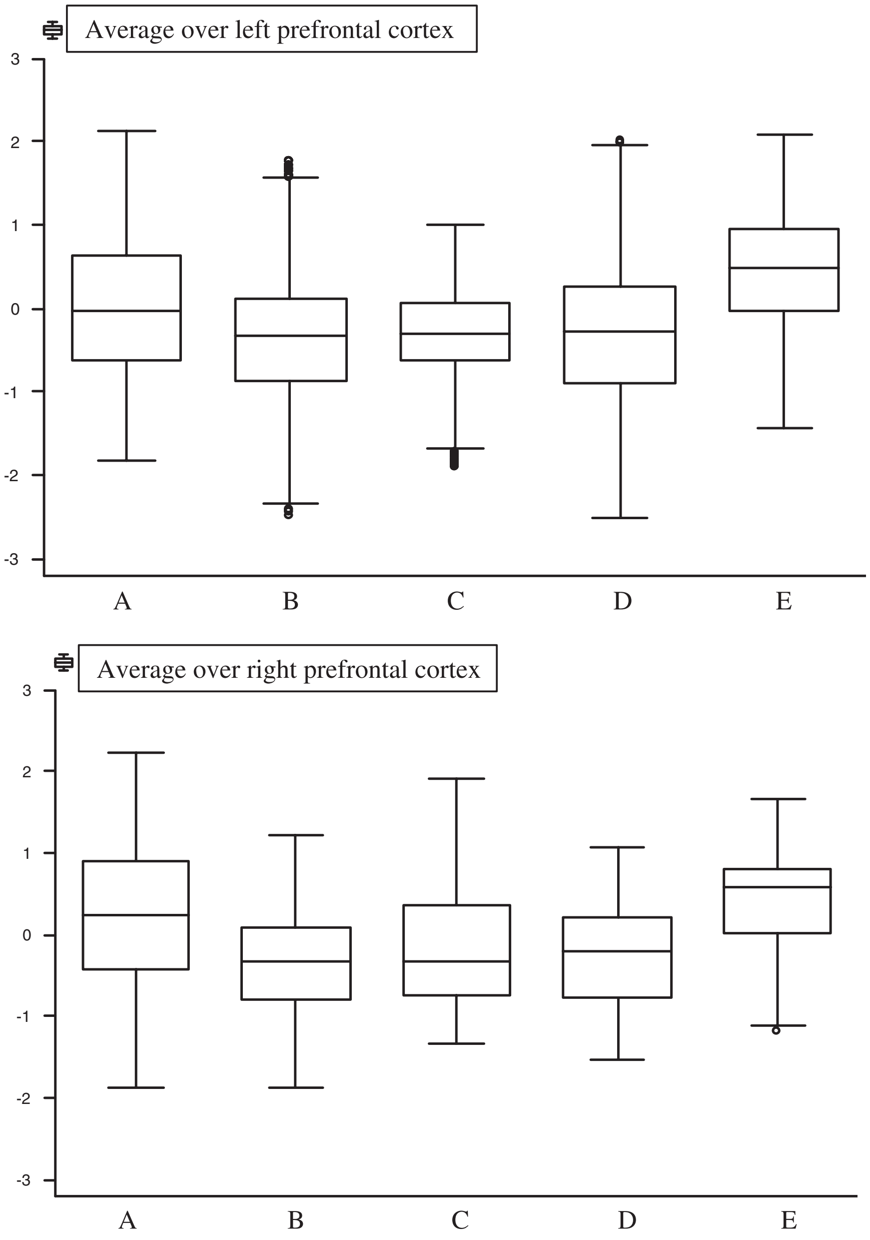

Figure 6. Left and Right Temporal Prefrontal Cortex Box and Whiskers Comparing All Manipulations.

the differences between the distributions are much more pronounced than between manipulations A and D.

Finally, Hypothesis 3 predicts that there will be no difference between manipulations D and E if subjects pay attention to their counterparts. This hypothesis can be rejected ($t = -15.91$, $df = 1{,}198$, $p < .001$ for the left hemisphere and $t = -20.55$, $df = 1{,}198$, $p < .001$ for the right hemisphere). These differences are rather dramatic and can be seen in the box and whiskers plots shown on Figure 6. The difficulty with calculating a mixed strategy in manipulation E may explain those differences. Subjects may simply be guessing a response, rather than accounting for their counterpart—hence the high positivity for manipulation E.

Discussion

The pattern of data is consistent with our conjectures about the role of a counterpart. First, we find that there are differences between the simplest decision problem and the simplest strategic problem. In both instances manipulations A and D are simple dominance solvable games. The solution for manipulation A is trivially easy, while the solution for manipulation D is a bit more difficult, but only because of the extra values on the screen. We find a large difference in the response time by subjects across these two manipulations. In part this is because subjects have to spend their time deciding what type of strategic game they face when in the fourth block of decisions and see both manipulation D and E decisions. What is interesting about these data is that in the interval from 500 to 800 milliseconds out from the keypress, the easier decision (A) has greater positivity than the more difficult decision (D). Standard parametric tests indicate that these two treatments are different and the direction of those differences is consistent with what we might expect if subjects are paying attention to a counterpart.

Second, we hypothesized that manipulations C and E would be similar, but they are not. Manipulation C involves a choice under uncertainty, in which nature makes a choice, but the probability is not revealed. Manipulation E is similarly designed, except that a human makes the choice rather than nature, but also with an unknown probability. If subjects did not take their counterpart into account these two manipulations are identical. The response times indicate that subjects spend much more time considering manipulation E. However, manipulation E also shows much higher positivity, something that is not consistent with the idea that greater pre-frontal cortex activity is being recorded. Hence, these results are mixed.

Third, in a simpler task, in which manipulations A and D look identical, except for the presence of a human counterpart, we find differences in pre-frontal cortex activity. Manipulation D has greater negativity, which is consistent with greater neuronal firing. However, we are troubled by the fact that a similar pattern does not hold for manipulation E in which a human counterpart makes a choice.

Conclusion

The motivation for this study is the idea that strategic interaction involves a "reading" of one's counterpart and that this reading may support the kind of "out of equilibrium strategies" that experimental game theorists often observe in the lab. This suggests, then, that the cognitive process by which people make decisions in situations that involve others should be different from the cognitive process by which they make (equally complex) decisions that do not involve a counterpart. One good example of the kind of difference we were looking for is between nonstrategic decisions made under uncertainty and strategic decisions that depend on the uncertain play of others. In the former, no "reading" of the other is possible, while in the latter such a "reading" may be critical to one's own choice. Our attempt to identify such differences in the electrical activity of the pre-frontal cortex produced mixed results. For the type of situation just described differences were observed, but for the other situation (the simple decision problem versus the dominant strategy game) we found no differences.

It would be unwise, however, to over interpret these preliminary results. There are a number of issues that can be improved in future work. First, it is clear that task complexity must be more carefully controlled for in comparing tasks involving a counterpart with those that do not. At the very least, we would need to compare tasks that generate similar distributions of response times and that present the same amount of information to subjects. Second, it would be useful to have more subjects. The extent of between subject variation in this experiment was greater than expected.

Beyond the conclusions of this specific study, we can also offer some general conclusions about the use of ERP methodology in experimental game theory. On the positive side, the ERP methodology is relatively noninvasive. Subjects can be fitted with the equipment 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 and with a temporal resolution that is much finer than PET or fMRI.

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.

ACKNOWLEDGMENT

This paper was prepared for the 2001 Annual Meeting of the Midwest Political Science Association. 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. Correspondence concerning this article should be addressed to Rick Wilson (Department of Political Science, MS 24, Rice University, Houston, TX 77251-1892). E-mail: rkw@rice.edu

REFERENCES

Bolton, Gary E., & Axel Ockenfels. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review, 90*(1), 166–193.

Camerer, Colin F. (2003). Behavioral game theory: Experiments in strategic interaction. Edited by C. F. Camerer and E. Fehr, *The roundtable series in behavioral economics*. New York: Russell Sage Foundation, Princeton University Press.

Cameron, Lisa A. (1999). Raising the stakes in the ultimatum game: Experimental evidence from Indonesia. *Economic Inquiry, 37*(1), 47–59.

Eckel, Catherine C., & Philip Grossman. (1996). Altruism in anonymous dictator games. *Games and Economic Behavior, 16*, 181–191.

Eckel, Catherine C., & Rick K. Wilson. (2005). Judging a book by its cover: Beauty and expectations in the trust game. *Political Research Quarterly* (forthcoming).

Fehr, Ernst, & 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.

Handy, Todd C., ed. (2004). *Event-related potentials: A methods handbook*. Cambridge: The MIT Press.

Holt, Charles A., & Jacob Goeree. (2001). Ten little treasures of game theory, and ten intuitive contradictions. *American Economic Review, 91*(5), 1402–1422.

McCabe, Kevin, Daniel Houser, Lee Ryan, Vernon L. Smith, & Theodore Trouard. (2001). A functional imaging study of cooperation in two-person reciprocal exchange. *Proceedings of the National Academy of Science, 98*(20), 11832–11835.

Nunez, Paul L., & Ramesh Srinvasan. (2005). *Electric fields of the brain: The neurophysics of EEG*. 2nd ed. Cambridge: Oxford University Press.

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

Pribram, Karl H., & Aleksandr R. Luria, eds. (1973). *Psychphysiology of the frontal lobes*. New York: Academic Press.

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*(1), 11–46.

Rugg, Michael D., & Michael G. H. Coles. (1997). *Electrophysiology of mind: Event-related potentials and cognition*. Cambridge: Oxford University Press.

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.

Stuss, Donald T., & Robert T. Knight, eds. (2002). *Principles of frontal lobe function*. Oxford: Oxford University Press.

Zani, Alberto, & Alice Proverbio. (2002). *The cognitive electrophysiology of mind and brain*. New York: Academic Press.