

## Inferior Parietal Lobule Supports Decision Making under Uncertainty in Humans

Timothy J. Vickery<sup>1</sup> and Yuhong V. Jiang<sup>1,2</sup>

<sup>1</sup>Harvard University, Department of Psychology, Cambridge, MA, USA and <sup>2</sup>University of Minnesota, Department of Psychology & Center for Cognitive Sciences, Minneapolis, MN, USA

**The optimal responses for many decisions faced by humans are ill defined, yet we are able to choose well by associating choices with outcomes, and employing this information in decision making. Previous studies suggest that the parietal cortex is involved in "uncertain" decision making, yet uncertainty is confounded with increased difficulty and attention. Here we aim to dissociate the role of parietal cortex in decision making and attention. Using functional magnetic resonance imaging we measured brain activity while participants played a "matching-pennies" game. We found that the inferior parietal lobule is involved in decision making under uncertainty, showing higher activity when the decision was uncertain rather than certain and when humans were given trial-by-trial feedback on choice outcomes than when they were not. Crucially, increasing attentional load with secondary tasks reduced inferior parietal activity when decisions were made under uncertainty, suggesting that general attention does not drive its activation. This pattern was consistent for visual or auditory feedback, and for direct (symbols representing wins and losses) or indirect (only the opponent's choices were shown) feedback. It contrasted with results from medial superior frontal gyrus, which was driven primarily by increased attentional load. We suggest that decision making under uncertainty is dissociable from general attention in the brain.**

**Keywords:** attention, decision making, inferior parietal lobule, reward

### Introduction

Decision making sometimes involves straightforward, objective criteria, such as determining the direction of a moving object, even if made in the presence of perceptual noise (Shadlen and Newsome 2001), whereas at other times it is inherently uncertain, as when we gamble on a roll of the dice, the weather on a particular day, or the behavior of another agent. Much research in decision making has confronted the former type, by having participants categorize, discriminate, or detect a signal among noise (Heekeren et al. 2004). Increasingly, researchers are exploring the neural machinery underlying uncertain decision making (Barraclough et al. 2004; Dorris and Glimcher 2004; Sugrue et al. 2004; Glimcher 2005). When choices are uncertain, humans must often take into account many sources of evidence, especially previous experience (Sutton and Barto 1998; Feltovich 2000; Camerer 2003). Thus, the historical association of rewards and punishments with choices becomes a key element in decision making. Here, we examine the role of inferior parietal lobule (IPL) in supporting decision making under uncertainty, as distinct from rule-based decision making, general attention, and executive functions. Decision making under uncertainty is construed broadly, to

include processes that compute action-reward associations, opponent and environment modeling, and the integration of these factors into a final choice.

Previous research has located several frontal and parietal regions that are more active during uncertain decision making than certain, rule-based decision making. Some regions, including the inferior parietal cortex, show activity that is graded in proportion to relative feedback about previous choice outcomes (Paulus et al. 2001; Nickerson 2002; Miller et al. 2005). For example, Huettel, Song, and McCarthy presented contexts that represented with varying degrees of certainty which of 2 responses were "correct." Blood oxygenation level-dependent (BOLD) responses in regions of parietal cortex (among others) were higher for less certain than more certain choices (Huettel et al. 2005). However, parietal cortex is sensitive to increased attentional effort and task demands under a broad range of cognitive tasks (Kanwisher and Wojciulik 2000; Shuman and Kanwisher 2004). Increased activation in the presence of uncertainty may simply reflect greater attentional demands under uncertain than certain conditions. At present, it is unknown whether parietal regions singled out for involvement in uncertain decision making show specialization for calculations that support optimal decisions of this type, or whether they simply reflect increased attentional demands (Bechara et al. 1998; Maunsell 2004).

Miller et al. (2005) attempted to dissociate decision making under uncertainty from attentional demands by subtracting activity in an *N*-back working memory task from those in an uncertain decision-making task. The subtraction left no activation in IPL, casting doubts on its specific role in decision making. However, this subtraction contrasted 2 very different tasks, in which a number of variables were altered, clouding the interpretation of this contrast.

In the present study we sought to clarify the role of IPL in decision making under uncertainty using functional magnetic resonance imaging (fMRI). We asked whether activity in IPL reflects increased attentional demands associated with increased decision uncertainty, or whether it reflects the involvement of this region in either calculating choices on the basis of feedback from prior trials, or tracking the associations between choices and feedback. We manipulated as directly as possible the overall effort required of the participants across different attentional conditions and under different feedback conditions. In our analysis, we contrast BOLD activation patterns in IPL to those in the medial superior frontal gyrus (mSFG), another region which has an indeterminate role in attention and decision making (Carter et al. 2000; Duncan and Owen 2000; Walton et al. 2003; Rushworth et al. 2004). These 2 regions provide us with a clear and compelling contrast.

A general goal of this research was to specify criteria answerable by fMRI that might be necessary, though perhaps not sufficient, to define candidate cortical regions responsible for making decisions on the basis of prior experience. Such a region would track information about one's own prior choices and their consequences (i.e., punishment and reward signals) in order to make the next decision, and/or integrate these signals to make a choice. We devised 5 criteria for identifying such a region: 1) its activity should be sensitive to observance of choice outcomes that form the basis for making an optimal decision. It should be more active when trial outcomes are revealed rather than concealed. This should be true on both a trial-by-trial basis, and over the course of blocks in which feedback is either given or withheld. 2) This sensitivity should be dissociable from general attentional effort. Any manipulation that increases attentional demands, but impairs decision making under uncertainty, should lead to *decreased* neural activity in this region. 3) It should show higher activity when decisions are made under uncertainty than when they are made under certainty, even if concrete rewards (e.g., financial rewards) are associated with both. 4) These regions should be sensitive to prior choice outcomes, whether or not the outcome is explicitly stated (e.g., a "win" or "loss" signal) or inferred by matching the participant's own choices with an opponent's choice. 5) The sensory modality in which prior choice outcomes are observed, and the general context in which the feedback is given (e.g., spatial or nonspatial context), should not matter.

Our criteria serve to dissociate decision making under uncertainty broadly from other important processes. For instance, we do not claim to differentiate processes that incorporate probabilistic action-reward associations from those that incorporate beliefs about opponents or environments. This coarser grain of analysis is an oft-neglected precursor to determining candidate regions for more specific functions in decision making. We speculate on the specific function of IPL in the general discussion.

The basic task used throughout this paper is the simple, repeated "matching-pennies" game with 2 players (e.g., Barraclough et al. 2004). One player is the "matching" player, whereas the other is the "mismatching" player. On each round, both players choose one of 2 options simultaneously. The matching player wins when the choices are the same, whereas the mismatching player wins when they differ. We framed this game, in the first experiment, as a "penalty-kick" game in which the participant was the goalie, and the computer-based opponent was the kicker. The kicker could choose only to kick left or right, whereas the goalie could choose to jump left or right. The goalie wins every time the kicker kicks in the same direction as the goalie jumps.

This task was chosen for its simplicity and the reliance of the optimal strategy upon past experience. When played repeatedly, the only way to gain advantage is to predict the opponent's next move, presumably by exploiting statistical dependencies in the opponent's choices. For example, if an opponent tends to choose left after choosing right, or to switch choices after losing, more games can be won by anticipating and countering this tendency. Against a more capable opponent, such as the computerized opponent employed in our studies (see "computer's algorithm" under methods), a player's best option is to eliminate statistical dependencies in his or her own responses. However, humans are generally

poor at producing "random" choice sequences (Baddeley 1966). The term "random" is ill-defined, a problem that contributes to difficulties in producing random sequences (Nickerson 2002). We use the terms random and unpredictable synonymously. Considered as a temporal sequence, "randomness" corresponds to the extent to which a pattern detection algorithm could not successfully predict any particular choice on the basis of preceding choices. To the extent that it is possible, behaving unpredictably is dependent upon both feedback and attention. For instance, participants are incidentally better at producing random sequences in a matching-pennies game context than they are at producing such sequences on demand (Rapoport and Budescu 1992). This probably arises, in part, from dynamic use of feedback information, because opponents can detect and counter statistical regularities. Indeed, feedback on performance can enhance randomization even in the absence of a game context (Neuringer 1986). Additionally, random sequence generation has been found to strongly depend on attention and working memory. When distracted by a secondary task load, such as holding a concurrent memory load, participants in random sequence generation tasks produce more stereotyped sequences (Towse and Valentine 1997; Baddeley et al. 1998).

The dual dependence of successful randomization on feedback and attention make it ideally suited to our purposes in this paper. Reliance on feedback implies that withholding feedback should make participants less capable of performing the task, and thus regions involved in computing choices based on prior choice outcomes should be less active. The reliance of random sequence generation on attention means that increasing cognitive load should interfere with this process, as well. Whereas regions involved in coordinating tasks and deploying attentional resources should show increased BOLD activation under high load, those that are involved in assisting choice-making based on prior choices and their outcomes should show reduced activation, similar to the consequences of withholding prior-trial outcome information.

This paper presents 2 fMRI studies that examined the feedback- and attention-dependent activation of IPL and other putative decision-making regions. The randomization statistics used to evaluate behavioral performance in this task require many trials to determine reliably, so a pilot study was conducted to confirm the feedback- and attention-dependent nature of the "matching-pennies" task. We then scanned participants in fMRI while they played this game. A random-effects analysis was conducted to locate regions of interest (ROIs) that responded to feedback, the attentional demands of the task, or an interaction of the 2 that would imply attention-dependent feedback processing. In experiment 2, a separate group of participants were scanned in a variant of this task. Using ROIs obtained from experiment 1, we examined the influence of feedback modality and feedback format.

## Materials and Methods

### *Behavioral Pilot Study*

The behavioral pilot study was designed to test the feedback- and attention-dependent nature of performance in matching pennies. From the above-cited literature, we expect enhanced performance when feedback is given and degraded performance when secondary tasks are required. Thus, participants played a matching-pennies style game with

or without trial-by-trial feedback, and under either low or high attentional load (with a demanding secondary tone-counting task applied under high attentional load).

### Participants

Twenty-seven individuals from the Harvard University community participated in this study for payment. The experiment lasted approximately 45 min, and participants were paid \$8, plus up to \$6 as rewards for winning.

### Procedure and Design

Participants played 6 “games” of repeated matching-pennies against a computerized opponent (described below). The task was framed as a “penalty-kick” game in which the participant played as the goalie, and the computer played as the kicker. Participants pressed left and right arrow keys to move the goalie. Each game consisted of 300 consecutive trials. After each move, a 500-ms feedback phase showed the results of the move. For every trial that the human goalie chose the same as the computer kicker’s choice, one point was awarded. Participants were told that the computer would attempt to counter patterns in responses, and that the computer would observe the outcomes of trials even when they themselves could not. During 3 of the 6 games, participants saw their opponent’s responses (*feedback* condition, in which the player was shown to either catch or miss the ball) and in the other 3 they did not (*no-feedback* condition, in which only the direction of the player goalie’s jump was shown). These 2 feedback conditions were crossed with 3 secondary task conditions. In the *no sound* condition, participants heard no sounds as they played the game; in the *ignore sound* condition they heard one tone played on each trial (low, 200 Hz; medium, 400 Hz; or high-frequency, 800 Hz), but were instructed to ignore them; and in the *attend sound* condition they heard the sounds and were required to count only the middle frequency tones. Block sequence was counterbalanced and randomized between participants.

Participants were penalized for slow responses (more than 2000 ms from trial start) and rapid key tapping (responding before the removal of feedback). Both penalties were signaled by a large “-1” that appeared in the place of the goalie. Detailed instructions and extensive practice with both tone counting and playing the game ensured that participants understood these rules fully before play began.

Participants received monetary reward by comparing their scores on each game to the mean score on single games of all participants who were tested before (pilot data for the first few participants). Participants won \$1 for each game in which they out-performed the mean. After *attend sound* blocks, they were also asked to enter the number of counted tones. They were automatically disqualified from

winning the money if their response in tone-counting differed from the true answer by 10 or more.

### Analysis

We examined the optimality of player responses using both their scores, and simple measures employed in prior studies of random sequence generation. Baddeley and others have studied this task as an “information generation” problem (Baddeley 1966). A sequence that is more stereotyped and patterned carries less information than one that is purely random. We employed a measure that uses adjacent response pairs (digrams). A sequence of length  $n$  contains  $n - 1$  digrams. A measure of the information in digrams is  $H_2$ , which is calculated as

$$H_2 = - \sum_{j \in \{L,R\}} \sum_{i \in \{L,R\}} (n_{ij} / (n-1)) \log_2 (n_{ij} / (n-1))$$

where  $n_{ij}$  is the number of occurrences of a particular digram ( $ij$ ) in a sequence (e.g., LL—2 left responses). We also applied an adjustment based on single-response frequency. Correcting for these frequencies employs the theoretical maximum information per response pair,  $2H_1$ , which is obtained using the following formula:

$$H_1 = - \sum_{j \in \{L,R\}} (n_j / n) \log_2 (n_j / n)$$

The overall measure we used is redundancy,  $R$ , where

$$R = 100[1 - H_2 / (2H_1)].$$

This measure is a percentage, where higher values represent greater redundancy (i.e., less random sequences). For details see Baddeley et al. (1998) or Attneave (1959). Though this statistic has its weaknesses, it is convenient because it provides a single, meaningful measure.

### Experiment 1

Our aim in this experiment was to locate regions of the brain that responded to feedback (i.e., trial-by-trial information about wins and losses), and to dissociate this response from attention-related activity. Experiment 1a consisted of a blocked design, whereas experiment 1b was an event-related fMRI experiment conducted during the same scanning session.

The design of experiment 1a was similar to the pilot study, without the *no sound* conditions. For convenience, the major characteristics of experiments 1 and 2 are summarized by Table 1. During a single scan, participants completed 4 games: 2 with feedback and 2 without feedback; 2 with a secondary tone-counting task load (*high load*) and 2 without counting (*low load*). Regions that track choice-reward relationships or determine choices based on those relationships should show increased activity when trial-by-trial feedback is available, but they should *not* be more active due simply to increases in the general attentional demands of a task.

**Table 1**  
Design parameters and conditions for all experiments

Experiment	1a	1b	2a	2b
Design	Blocked	Event-related	Blocked	Blocked
Game type	Penalty kick	Penalty kick	Matching game	Matching game
Feedback type and symbols	Visual. +/- Symbols for wins and losses.  * For no-feedback.	Visual. +/- Symbols for wins and losses.  * For no-feedback.	Visual. Computer’s choices were signaled by a small square that stood for choice A and a large square that stood for choice B. A union of the small and large square represented no-feedback trial outcomes.	Auditory. Computer’s choices were signaled by a spoken “A” or “B.” A spoken “C” represented no-feedback trial outcomes.
Secondary task	Auditory. Count high-frequency tones.	None.	Auditory. Count number of spoken “A” among “B.”	Visual. Count number of visually presented large among small squares.
List of conditions	Feedback/low load Feedback/high load No-feedback/low load No-feedback/high load	Feedback No-feedback	Feedback/low load Feedback/high load No-feedback/low load No-feedback/high load Rule-based/low load Rule-based/high load	Feedback/low load Feedback/high load No-feedback/low load No-feedback/high load Rule-based/low load Rule-based/high load

## Participants

Eighteen participants (8 female) recruited from the same general population completed experiment 1 for financial compensation. One subject's data were discarded due to excessive head motion.

## Experiment 1a: Procedure and Design

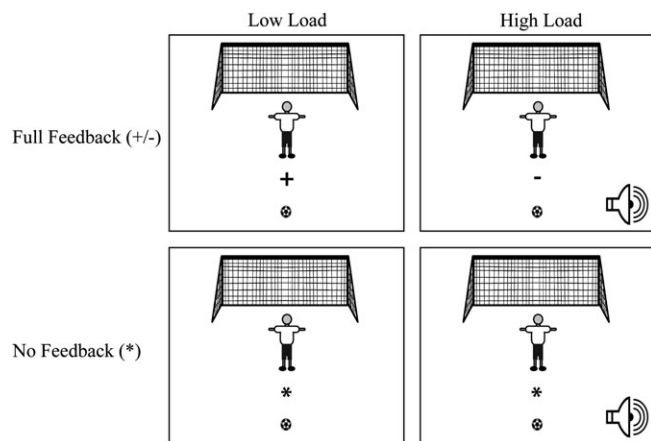
Participants played the penalty-kick game against the computer (Fig. 1; Table 1), controlled by the algorithm used in the behavioral pilot study (described in "Computer's algorithm" below). In this experiment, the human played as "goalie" and won whenever the choices "matched."

Experiment 1a was a  $2 \times 2$  design, with 2 feedback conditions crossed with 2 attentional load conditions. During each block, outcomes were either withheld or provided. When outcomes were withheld, the participants only saw the fixation point turn into an asterisk during the feedback phase of every trial. When outcomes were provided, the fixation point turned into either a plus or a minus sign for wins and losses, respectively. When no response was received, the fixation point turned into the characters "-1?" to signal a missed response, and a point was taken off their score.

To manipulate attentional load, we required a secondary task to be completed. During each trial, either a high- or low-frequency tone played through the scanner-room speaker. Under *high-load*, participants counted the number of high-frequency tones that played during that block. At the end of each block, 4 consecutive numbers appeared for 4 s, with one matching the actual number of target tones played, and participants were asked to select the target number. Under *low-load*, participants ignored the tones, but they selected a high-lighted (green among red) number during the number response phase. In both cases, if the correct response was not made, the subject was not eligible for reward regardless of the penalty-kick performance.

Blocks consisted of 24 consecutive trials under identical experimental conditions. Trials were 2-s long and consisted of a 1500-ms choice phase followed by a 500-ms feedback phase. All 4 conditions were present in every scan, in a different order. Blocks were preceded by a 16-s fixation period, and a 4-s instruction phase with a short description of their task, and they were followed by the 4-s secondary task response phase and a 16-s fixation period. In total, each scan of experiment 1a lasted 304 s. Participants completed between 4 and 8 scans of this task, as time allowed.

Participants were rewarded if their scores in a single block (or "game") were higher than the average block scores of prior participants (and if they selected the correct number at the end of that block). Each winning game resulted in a \$0.50 reward for experiment 1a.



**Figure 1.** Illustration of task in experiment 1. Participants had 1500 ms to direct their character to "jump" left or right, whereas the computer simultaneously chose to "kick" left or right. This was followed by 500 ms of feedback. In some blocks, participants received trial-by-trial feedback in the form of "+" and "-" symbols that represented wins and losses, whereas on other blocks an asterisk was presented as feedback regardless of the outcome. During feedback, a tone played (high or low frequency). When asked to count tones in some blocks, the participants had to count high-frequency tones and report the number at the end of the block

## Experiment 1b: Procedure and Design

Experiment 1b used an event-related design to ensure that feedback-related activity observed in experiment 1a was not due solely to state-based activation. For example, feedback could trigger overall strategic differences between *feedback* and *no-feedback* blocks. Feedback incorporated into uncertain decision making should modulate brain activity on a trial-by-trial basis, as well.

Participants played the same game, but with only 50 *feedback* and 50 *no-feedback* trials, which were exactly like individual trials from experiment 1a, but without sounds and no secondary task. These trials were intermixed with 50 fixation trials in which only the fixation marker appeared. Conditions were presented in a pseudorandom order determined by the optseq2 algorithm to maximize the design's efficiency (Dale 1999; <http://surfer.nmr.mgh.harvard.edu/optseq/>). Participants completed between 4 and 8 scans of this experiment. They were paid \$2.00 for every scan on which their score beat the average of prior participants' scores.

## Experiment 2

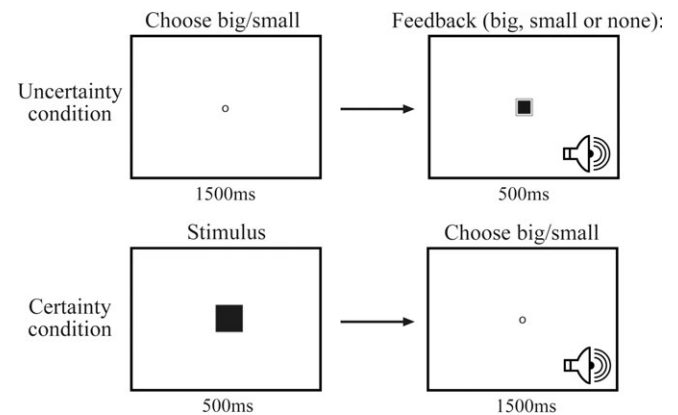
Experiment 2 was conducted to replicate and extend the results of experiment 1 in a new group of participants. The task and logic of the design were similar, but rule out a number of factors that could explain experiment 1's results (spatial processing and the type or modality of feedback). Further, we included a certain, rule-based decision-making task that was also associated with financial rewards, in order to verify that feedback-related activity was specific to the uncertain decisions required of our game task.

## Participants

Twelve individuals (4 female) participated in experiment 2.

## Procedure and Design

Experiment 2 employed a similar design as experiment 1a, differing only in the following ways (see Fig. 2 and Table 1). First, the game was not framed as a "penalty-kick" game; instead, it was described as a "matching game" in which the participant and computer both chose from 2 possibilities (choices "A" and "B"). The participants were told that they won when the choices matched, whereas the computer won when the choices mismatched.



**Figure 2.** In experiment 2, participants played a more abstract version of the same game. In the "visual" game (shown here), they chose small or large squares as the computer simultaneously made a choice between small and large. Humans won trials in which choices matched, and the computer won if the choices mismatched. As before, some blocks gave full feedback (revealing the computer's choice after every trial), whereas other blocks withheld feedback (revealing a combination of both responses as a placeholder), and participants were also sometimes required to count the number of times a voice said "B" rather than "A." On other scans, the counting and feedback modalities were reversed (the auditory game). Under uncertain-decision conditions, the game was played as in experiment 1. Under certain-decision conditions, the computer's choice appeared first, and participants simply had to select the same choice to win for that trial.

Secondly, an additional factor of feedback modality was introduced, split into 2 separate sets of scans reported separately as experiments 2a and 2b. Feedback came either in the visual domain (experiment 2a), in which a small or large square represented the computer's choices (choices A and B, respectively), or in the auditory domain (experiment 2b), in which a voice speaking the letter "A" or the letter "B" represented the computer's choices. When outcomes were not provided, the participants saw or heard uninformative feedback (corresponding to the asterisks of experiment 1): a union of the small and large squares (2a), or a vocalized "C" (2b). In experiment 2a, the high-load secondary task was performed by counting spoken "B" among "A" sounds, and in experiment 2b the high-load secondary task required counting visually presented large squares among small squares. Thus, the modality of the primary game task feedback and secondary counting task stimuli was reversed between experiments 2a and 2b.

Experiment 2 also included a "certain" feedback condition played under both *high load* and *low load*. In the certain task, trials began with a 500-ms period in which the computer's choice was shown or played through the scanner speaker, followed by a 1500-ms response period. The participant only had to make the same choice as the computer under these conditions. Participants were still rewarded as long as their error rate was lower than 12.5% and that they performed correctly on the secondary counting task.

Experiments 2a and 2b (4 scans each) were completed in separate chunks within the same session. Each scan contained 6 blocks, with the load manipulation (*low load* or *high load*) crossed with the game feedback manipulation (*feedback*, *no feedback*, and *certain* decision-making tasks). Order was counterbalanced across the 4 scans, and the order of the 4 scans was randomly assigned between participants. Rewards were increased to \$1.00 per block, otherwise determined as in experiment 1a.

#### Computer's Algorithm

The algorithm used in all experiments is adapted from recent studies by Daeyeol Lee and colleagues (Barraclough et al. 2004; Lee et al. 2004). The computer opponent consisted of an algorithm that tracked all subject responses from the start of a scan (or "block" in the behavioral pilot experiment). All games were the binary choice matching-pennies game, so choices will be referred to as A and B. On each trial, the sequence was examined for local sequential dependencies for sequences of up to 4 consecutive responses. For example, if the last 4 plays were ABAB, the computer examined the choice history for bias, in the probability of an A response  $p(A)$ , bias after choosing B,  $p(A|B)$ , bias after choosing A then B,  $p(A|AB)$ , and so on. These values were tested for deviations from unbiased choices using the binomial test. The computer used the test with the lowest  $P$  value and biased its own response to exploit this tendency. For example, if the worst violation was  $p(A|AB) = 0.7$ , then the computer would choose B with a probability of 0.7 (if the computer was mismatching). If no test violated unbiased expectations, then the computer chose randomly.

#### Functional Imaging Methods

Participants were scanned in a Siemens Allegra 3T Trio (Erlangen, Germany) scanner at the Martinos Center for Biomedical Imaging (Charlestown, MA). Participants viewed displays projected onto a screen by an LCD projector through a mirror attached to the head-coil. A computer running MATLAB with Psychophysics Toolbox extensions (Brainard 1997; Pelli 1997) was used to present stimuli and collect behavioral responses.

We first acquired high-resolution anatomical images using standard protocols. A standard  $T_2^*$ -weighted gradient echo pulse sequence (time echo 30 ms, flip angle 90°) with time repetition (TR) of 2000 ms was used in all functional runs. Twenty-eight 4-mm-thick (3.13 mm × 3.13 mm in-plane, 0 mm skip) slices parallel to the anterior commissure-posterior commissure (AC-PC) line were collected for each functional scan.

In experiment 1, scans were bunched, such that the 28 slices were obtained within the first 1568 ms of the 2000-ms TR, leaving a 432-ms silent period for the presentation of auditory sounds, and the tones were played in the silent interval in order to facilitate auditory

detection. In experiment 2, scans were not bunched, but before collecting functional data each subject verified that they could clearly discriminate the letters spoken through the scanner speaker over scanner noise and the ear protection. All participants reported that they could clearly hear the tones in experiment 1 and the spoken letters in experiment 2.

#### fMRI Analysis Methods

Data were analyzed using SPM99 (<http://www.fil.ion.ucl.ac.uk/spm/software/spm99/>) and in-house software for MATLAB. All data were preprocessed using this software, including standard realignment, motion correction, and spatial smoothing procedures. For experiment 1a, a whole-brain random-effects group analysis was carried out for the main effect of knowledge of prior choice outcomes (feedback), the main effect of attentional load, and their interaction ( $P < 0.001$ ). This analysis yielded 16 ROIs. These anatomical coordinates were used to define ROIs for each subject's data (radius = 9 mm) for each of the other experiments. Within each ROI, percent signal change relative to the fixation baseline was averaged for each subject and condition. All analyses treated participants as random factors. For experiment 1b, time course data for each subject and condition was extracted, and then peak responses relative to baseline were extracted from each time course.

## Results

### Behavioral Pilot Study

Results of the pilot study confirmed the dependence of performance on the feedback and attentional resources. The computer opponent we employed was successful at defeating human players. Humans won approximately 42.4% of trials, which significantly differed from the expected proportion of 50% if the computer was making choices indiscriminately. The computer was effective at defeating the humans under all conditions (one-sample  $t$ -tests, all  $P < 0.001$ ). Performance in *no sound* was indistinguishable from *ignore sound* conditions, so data was combined for these 2 conditions (henceforth the *low load* conditions). We found that performance was better with feedback than without (44.6% vs. 40.3%;  $F_{1,26} = 11.4$ ,  $P < 0.005$ ). Participants also had lower scores under *high load* than *low load* (41.3% vs. 43.6%;  $F_{1,26} = 4.11$ ,  $P = 0.053$ ). There was no significant interaction ( $F < 1$ ).

We also examined the optimality of player responses using percent redundancy, as described in the methods section. Participants showed a main effect of attention, such that percent redundancy scores were lower (and thus, responses were more random) under *low-load* than *high-load* conditions (1.20% vs. 2.28%,  $F_{1,26} = 55.64$ ,  $P = 0.025$ ). The main effect of feedback on percent redundancy showed a trend towards lower percent redundancy when feedback was given than when it was withheld (1.33% vs. 2.15%,  $F_{1,26} = 2.72$ ,  $P = 0.11$ ). Post hoc  $t$ -tests showed significantly lower percent redundancy in the *feedback, low load* condition (0.88%) than the *no-feedback, low load* condition (1.51%) [ $t^{26} = 2.33$ ,  $P = 0.028$ ], but an insignificant difference between *feedback, high load* and *no-feedback, high-load* conditions [ $t^{26} = 1.13$ ,  $P = 0.27$ ]. The interaction between and feedback conditions and secondary task on percent redundancy was not statistically significant ( $F < 1$ ).

Taken together, these results imply that performance on this task depends on feedback and attention. Removing feedback and increasing attention load had the same general effect of producing lower scores against the computer and higher percent redundancy in response sequences. Our fMRI

experiments exploit the feedback- and attention-dependent nature of this task, and ask “what regions of the brain are involved in making decisions under uncertainty?” These regions should show activation in response to feedback, and interference from secondary tasks.

### Experiment 1a

For experiment 1a, we conducted a random-effects group-based analysis to locate brain regions that were significantly modulated by feedback (*feedback > no feedback*), attention (*high load > low load*), and their interaction (greater feedback modulation under *low load* than *high load*). Two of the resulting ROIs were in the parietal lobe (both in right IPL [rIPL]), and both showed similar patterns of activation. One ROI was in the right caudate, one was located in the thalamus, 2 were in the temporal lobes, and the remaining ROIs were located in the frontal cortex (bilateral insula and medial frontal regions, as well as middle frontal gyrus). These regions are listed, with Montreal Neurological Institute (MNI) coordinates, in Table 2.

For this study, we focused on 2 ROIs that show strikingly different profiles of activation. mSFG (MNI: [-12, 9, 51]), in the vicinity of supplementary motor area and anterior cingulate, showed greater activity under high-load than low-load conditions, both when participants received prior-trial outcomes [ $t^{16} = 4.85, P < 0.001$ ] and when they did not [ $t^{16} = 8.19, P < 0.001$ ] (Fig. 3b). mSFG activity was not modulated by whether participants received feedback on prior-trial outcomes ( $F < 1$ ), suggesting that mSFG was sensitive to attentional effort involved in a task, but not to how decisions were made. However, the interaction between feedback and attentional load was significant ( $F_{1,16} = 9.35, P = 0.008$ ), due to slightly higher activity for feedback than no-feedback trials under low load [ $t^{16} = 2.20, P = 0.043$ ] and slight lower activity for feedback than no-feedback trials under high load [ $t^{16} = 2.15, P = 0.048$ ]. Overall, the profile of activity for this region implies a primary involvement in attention, rather than feedback processing.

In contrast to mSFG, the rIPL (MNI: [48, -48, 54]) was modulated by feedback (higher activity with feedback than without it) and showed reduced activity under high load compared to low load, suggesting that its activation was not driven by increased attention. Under low-load, rIPL activity was significantly higher when participants received prior-trial

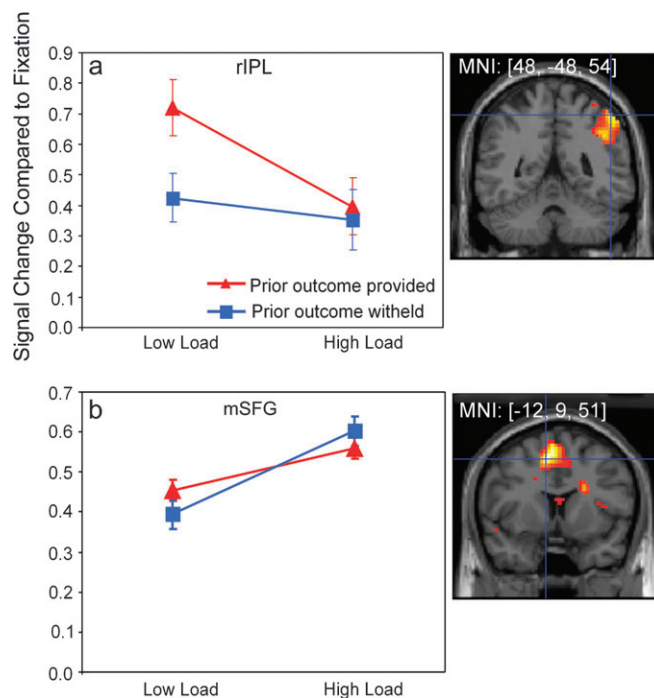
outcomes than when they did not [ $t^{16} = 5.89, P < 0.001$ ] (Fig. 3a). This increase was not produced by greater task difficulty involved in processing prior-trial outcomes, because participants response times were not significantly different when prior-trial outcome was provided (768 ms) than withheld (775 ms,  $t < 1$ ). Attentional load did modulate rIPL activity but in a direction opposite to its effects on mSFG. rIPL activity was reduced under high-load conditions when participants received prior-trial outcomes [ $t^{16} = 4.79, P < 0.001$ ], but was not modulated by attention when feedback was withheld ( $t < 1$ ), reflecting a highly significant interaction between attentional load and feedback ( $F_{1,26} = 53.6, P < 0.001$ ). The significant interaction between load and prior-trial outcomes indicates that the ability to track and make choices based upon previous choice-reward relationships was impaired under high attentional load, suppressing rIPL activity. All interactions involving regions (rIPL vs. mSFG) and either attentional load, feedback, or both were significant (all  $P < 0.05$ ).

We examined the relationship between behavior and rIPL activity by calculating the correlation between each participant's average score and the responsiveness of rIPL to feedback (subtracting mean BOLD signal under high load from activity under low load, when feedback was provided). No significant correlation was observed ( $r = 0.30, P > 0.20$ ). This lack of correlation may be due to unreliability of the individual difference measures. Because each block only consisted of 24 trials, individual differences did not have an opportunity to be expressed in the score measure, and variance in the scores

**Table 2**

ROIs obtained from whole-brain analysis of experiment 1 results

Region (BA)	MNI coordinates (x, y, z)	Contrast used to obtain
R. mSFG (BA 32)	[-12 9 51]	Attention (high > low load)
R. caudate	[18 0 24]	Attention (high > low load)
R. insula (BA 13)	[39 12 9]	Attention (high > low load)
L. insula (BA 13)	[-33 24 3]	Attention (high > low load)
R. middle temporal gyrus (BA 21)	[54 30 0]	Attention (high > low load)
R. medial frontal gyrus (BA 6)	[3 42 39]	Feedback > no-feedback
rIPL (BA 40)	[51 -51 36]	Feedback > no-feedback
R. middle frontal gyrus (BA 6)	[39 0 45]	Feedback > no-feedback
R. SFG (BA 9)	[39 39 33]	Interaction
Thalamus	[9 -9 3]	Interaction
rIPL (BA 40)	[48 -48 54]	Interaction
R. middle temporal gyrus (BA 21)	[63 -30 -12]	Interaction
R. SFG (BA 10)	[15 57 3]	Interaction
R. insula (BA 47)	[36 24 -6]	Interaction
L. insula (BA 13)	[-33 15 -6]	Interaction
R. medial frontal gyrus (BA 6)	[3 30 39]	Interaction



**Figure 3.** Percent signal change over fixation for experiment 1 in (a) rIPL and (b) mSFG as a function of feedback and attention manipulations ( $N = 17$ , mean  $\pm$  between-subjects SEM). rIPL activity was higher when feedback was provided than when it was withheld, and activation reduced significantly when a secondary task was required. In contrast, activity in medial superior frontal gyrus (mSFG) did not depend on feedback, but increased in response to attentional demands (rIPL, MNI [48, -48, 54]; mSFG, MNI [-12, 9, 51]). Brain images shown are from the whole-brain analysis,  $P < 0.001$ . Graphs plot activity in a 9-mm ROI surrounding the pinpointed region

was inherently random. A similar analysis of experiment 2's data was also unfruitful, subject to the same issues. For instance, scores between the visual (2a) and auditory (2b) games were poorly correlated within the same participants.

### Experiment 1b

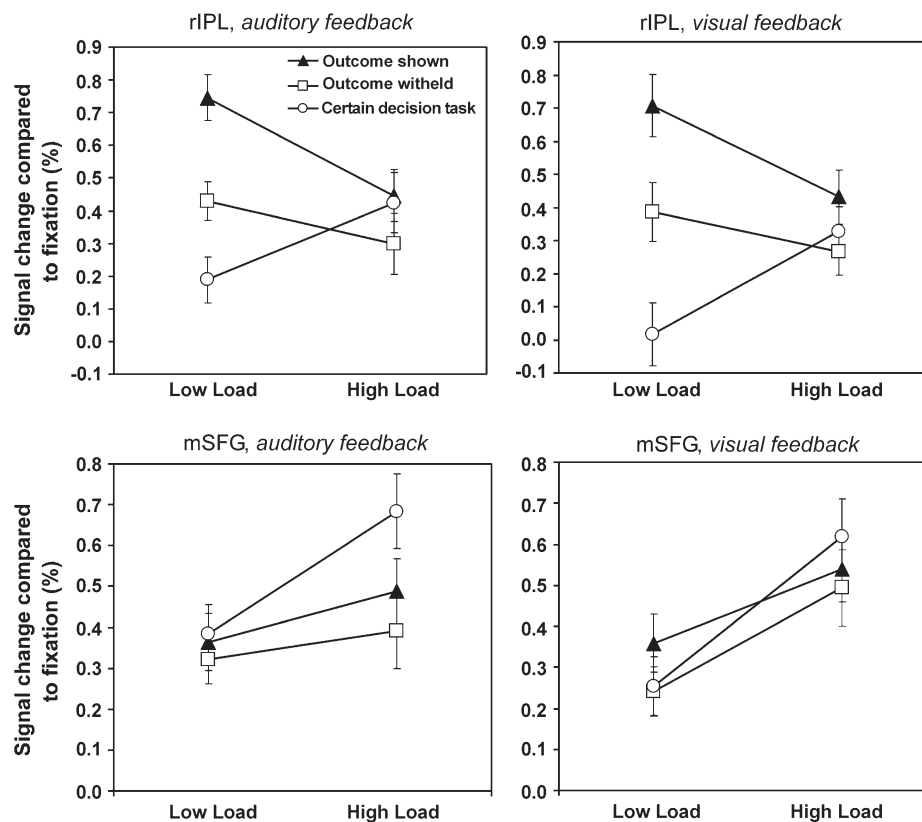
The event-related experiment 1b showed that rIPL (but not mSFG) was modulated by feedback on a trial-by-trial basis. Experiment 1b included only *feedback* and *no-feedback* conditions, randomly intermixed. We found peak activation in averaged BOLD time courses in response to *no-feedback* trials and *feedback* trials, separately analyzing *wins* and *losses* for the *feedback* condition. The resulting analysis showed a significant trial-by-trial modulation of activity in rIPL ( $F_{2,32} = 4.75$ ,  $P = 0.016$ ). A post hoc analysis showed higher peak BOLD activation for averaged activity in response to *wins* and *losses* over *no-feedback* trials [PSC: 0.081% vs. 0.024%;  $t^{16} = 6.66$ ,  $P = 0.02$ ], and marginally higher responses to *wins* than *losses* [PSC: 0.081% vs. 0.044%;  $t^{16} = 1.84$ ,  $P = 0.084$ ]. In contrast, mSFG was not significantly modulated by trial-by-trial feedback ( $F_{2,32} = 1.42$ ,  $P = 0.26$ ; PSC of 0.055%, 0.026%, and 0.030%, for wins, losses, and no-feedback, respectively).

### Experiment 2

Experiment 2 serves as an important replication and extension of the results of experiment 1a. Experiment 2 examines activity at the same ROI coordinates defined by Experiment 1 in an independent dataset. In experiments 2a and 2b ( $N = 12$ ), the

game was explained as a standard "matching-pennies" task that involved nonspatial choices. By pressing one of 2 keys, the player tried to match the computer's choice of either a spoken "A" or "B" (auditory task) or a centrally presented small or big square (visual task). Although the soccer penalty-kick task may have induced spatial imagery of the ball moving to one side of the goal, the new tasks were designed to induce minimal spatial processing. In addition, we provided feedback not by directly informing participants whether they won or lost, but by revealing their opponent's choice upon which the outcome could be inferred.

A ROIs analysis focusing on rIPL (functional coordinates defined by experiment 1, MNI: [48, -48, 54]) revealed that, even under these conditions, rIPL showed greater activation when prior choice outcomes were revealed than when they were concealed (auditory:  $F_{1,11} = 50.3$ ,  $P < 0.001$ , visual:  $F_{1,11} = 23.3$ ,  $P < 0.005$ ) (Fig. 4a,b). Once again, rIPL activity did not increase under high-load regardless of whether feedback was given, arguing against its role in general attention. Instead, rIPL activity was suppressed by high attentional load, particularly when participants received prior choice outcomes (auditory:  $F_{1,11} = 19.2$ ,  $P < 0.001$ , visual:  $F_{1,11} = 9.13$ ,  $P < 0.01$ ). This pattern of activation was virtually identical when participants played a visual game (match small and big squares while counting spoken "B") and when they played an auditory game (match spoken "A" and "B" while counting big squares),  $P > 0.4$  for all modality interaction effects. Thus, rIPL was sensitive to whether a decision was made on the basis of prior choice



**Figure 4.** Percent signal change over fixation for experiment 2 in (a) rIPL and (b) medial superior frontal gyrus (mSFG) as a function of feedback and attention manipulations ( $N = 12$ , mean  $\pm$  between-subjects SEM). rIPL activity was again higher when feedback was provided than when it was withheld, and activation reduced significantly when a secondary task was required. In contrast, activity in mSFG did not depend on feedback, but increased in response to attentional demands, as before. Activity in response to the certain decision making task was lower than uncertain with or without feedback for rIPL, and activity under these conditions increased with a secondary task (rIPL; mSFG).

outcomes, independently of whether the decision-making task involved spatial or nonspatial choices, whether the feedback was delivered through visual or auditory modality, and whether the feedback was explicitly stated or inferred.

Critically, the pattern of results shown above applied only when participants made decisions under uncertainty. When participants chose after seeing the computer's choice, there was no ambiguity and the human player won on 94% of trials. Even though participants knew the outcome of their choice on every trial, activity in rIPL was much lower in this condition than when decisions were made under uncertainty under low load (all  $P < 0.01$ , both modalities). Thus, knowledge of prior choice outcomes is not what drives rIPL activity. Only when the information was incorporated in future decisions did rIPL increase its activation. Finally, increasing attentional load when decisions were made under certainty increased rIPL activity [auditory:  $t^{11} = 3.91$ ,  $P < 0.005$ , visual:  $t^{11} = 4.45$ ,  $P < 0.001$ ], in sharp contrast to the pattern seen under uncertain decision making.

Analysis of mSFG again revealed its involvement in general attention (Fig. 4c,d). mSFG activity was not consistently higher when decisions were made under uncertainty than certainty, suggesting that it is not specifically involved in decision making under uncertainty. Furthermore, mSFG activity was greater at high-load than low-load (auditory:  $F_{1,11} = 10.4$ ,  $P < 0.001$ ; visual:  $F_{1,11} = 24.0$ ,  $P < 0.001$ ), even when decisions were uncertain (both  $P < 0.001$ ). The effect of decision certainty and the effect of attentional load both interacted significantly with brain region (all interactions between brain region, attentional load, and condition were significant,  $P < 0.001$ ).

#### Other ROIs

Of the ROIs derived from experiment 1a (see Table 2), only rIPL (both ROIs), medial frontal gyrus (mFG, anterior to mSFG ROI discussed above, MNI: [3, 30, 39], Brodmann area [BA] 6), and right SFG (rSFG, [39, 39, 33], BA9) satisfied all of the 5 criteria laid out in the introduction. All of these showed very similar patterns of activation to rIPL when analyzed across all experiments.

#### Discussion

In this study, we started with an extremely simple game played under uncertainty. Optimal behavior was defined as randomization over choices. Our behavioral pilot study pitted participants against a computerized algorithm that was designed to exploit statistical dependencies in their data. We found that participants played strategies that were closer to optimal when they were allowed to see the outcome of every trial, and when they played the game under low attentional load. These characteristics were exploited in 2 fMRI experiments that were designed to examine regions involved in decision making under uncertainty. In particular, we were interested in those regions that putatively support the connection of prior-trial outcomes to subsequent choice. We showed that fMRI activity in rIPL was consistent with an area that either incorporates feedback information on a trial-by-trial basis to compute a choice, or processes reward-related signals under uncertainty. Activity in rIPL was remarkably consistent across 2 experiments that tested separate groups of participants under very different conditions of stimulation.

In the introduction, we proposed a set of characteristics that a brain region involved in decision making under uncertainty

should manifest, 2 of which were confirmed in rIPL in experiment 1, namely, 1) its activity is sensitive to prior choice outcomes which form the basis for making an optimal decision, and 2) this sensitivity is dissociable from general attentional effort. Activity in rIPL was highest when trial-by-trial outcomes were shown under *low load* conditions. Under *high load*, however, activation was greatly suppressed. In contrast, activity in mSFG increased under *high load* and was modulated only slightly by feedback, implying that this region is attention-dependent. In addition, rIPL was active in response to feedback over blocks of several trials, and over individual trials in an event-related design, compared to blocks and trials in which no feedback was provided.

Experiment 2, analyzed using the anatomically defined ROIs derived from experiment 1, gave statistically independent verification of our results in rIPL and mSFG. It also allowed us to examine yet other criteria that a brain region involved in decision making under uncertainty should manifest. First, we stated that this region should be more involved in decision making under uncertainty than in decision making under certainty (criterion 3). This was shown in rIPL, where certain decision-making tasks elicited lower activation under *low load* than uncertain decision-making tasks, even though financial reward was associated with accurately performing the task.

We also suggested that 4) such a region should be sensitive to prior choice outcomes whether the outcome is explicitly stated or inferred, and whether it is delivered through visual or auditory modality. rIPL was clearly modulated by feedback in experiment 2, even though feedback came only in the form of an image or sound that indicated the computer's choice. In order to successfully infer the outcome, this symbol had to be translated into a choice and compared with the participants' own choices. Despite the removal of explicit valence from the feedback, rIPL showed nearly an identical pattern of activation in experiments 1 and 2.

Our final criterion was that 5) its pattern of activity should generalize to feedback in both modalities, and to nonspatial as well as to spatial contexts. The pattern of activity in rIPL did not depend on modality of the feedback: the same basic pattern of results was found for feedback in either domain. Concerns arose after the first experiment that because our choices were framed as "left" and "right" kicks in a penalty-kick game, activity in rIPL could be the result of spatial imagery or reasoning, a role that portions of parietal cortex are thought to play (Chafee et al. 2007). In experiment 2, we framed the game as a simple matching-pennies task, a context that should reduce the possibility of introducing spatial components to the task, but this seems to have had little effect on rIPL.

Although parietal cortex has been implicated in human decision making, much of previous research was plagued by confounding factors, especially given the role that parietal cortex is thought to play in attention (Kanwisher and Wojciulik 2000; Maunsell 2004). By manipulating attentional load directly and dissociating the role of parietal cortex in decision making from attentional effort, the experiments presented here have demonstrated for the first time that the IPL has specific functions in human decision making, which are dissociable from its functions in attention. We showed that activity in rIPL increases when a decision is made under uncertainty, and when a decision is made on the basis of prior choice outcomes. This pattern was observed whether the decision task involved visual or auditory modality, demonstrating that the parietal

involvement in decision making is modality-independent. Computation in rIPL is modulated by attentional load because a secondary task interferes with processing of prior choice outcomes, yet such attention-modulation is dissociable from attentional enhancement under high-load, as is shown in mSFG. The different response patterns in rIPL and mSFG to prior choice outcomes and attentional load suggest that decision making under uncertainty is dissociable from general attention in the human brain.

What is the specific role played by IPL in decision making under uncertainty? Here, we defined decision making under uncertainty quite broadly, and included in this definition a number of possibly separable subprocesses. First, a necessary precursor for obtaining our results is an association between action and outcome. Participants must have been tracking the relationship between their choices and the computer opponent's choices in order to improve performance when feedback was given. Secondly, there must be a way to integrate this information and make a choice. That is, some process must take the probabilistic associations and make a decision on their basis. Finally, there are a number of other candidate processes that can contribute to a decision, as well. For instance, the participants may have constructed an internal model of opponents and/or the environment, which would be continuously updated by outcomes that confirmed or violated the "theory" of the participant. This notion is similar to beliefs or models within model-based reinforcement learning algorithms (e.g., Fudenberg and Levine 1998).

Our data support a role for rIPL in one of these subcomponents, as opposed to any general attentional function. We speculatively suggest that the most likely candidate process implemented by rIPL is the integration of reward and other factors into either a value representation corresponding to each choice, or into choice-making itself. One reason for this conclusion is that, even when feedback was withheld, participants showed greater activity in rIPL than when the choice was certain or rule based. Thus, activity was not completely driven by external rewards and punishments. We reason that rIPL activity in our experiments does not likely reflect the updating of internal models or beliefs, on the basis that our task was extremely simple. Given that plenty of practice was obtained prior to scanning, the relatively few likely candidate strategies, and the fact that the opponent was static in its own strategy, we doubt that opponent modeling (if it exists) would have been so active as to drive neural activity in this region so consistently and strongly.

Could strategic exploration explain our results? Daw and colleagues have shown that a region near the intraparietal sulcus shows increased activity in response to exploratory choices (Daw et al. 2006). Our results rule out a simple switch-or-stay explanation for our data: the certain choice condition of experiment 2 required a number of switch versus stay choices that was statistically indistinguishable from the number of switches elicited by our uncertain choice conditions. It is possible that participants were changing between global strategies (e.g., win-stay/lose switch and win-switch/lose stay) on the basis of feedback, activating rIPL, but we view this as unlikely. First of all, our ROI in parietal cortex was dorsal and posterior to the locations in parietal cortex probed by Daw et al. Secondly, strategic shifts would only be expected to occur in this task on a relatively small number of trials. Even though our blocked design contained short, 24-trial blocks, strong

reward-related activity was elicited. Further, a robust event-related signal in response to reward was also observed, even though exploratory strategic shifts would not be expected to occur often.

Further research should probe the relationships between rIPL and the other regions implicated by this study for their role in uncertain decision making. For example, it is possible that these regions are performing entirely different computations (like those delineated above) that are both feedback and attention dependent; for example, one region may integrate many factors to calculate an expected utility value for choice, whereas another may detect exploitable patterns in opponent responses. On the other hand, they may be closely related, interacting to complete common goals in decision making under uncertainty. Other questions remain about the roles of these regions in distinct types of uncertainty. Here, uncertainty arose due to ambiguity about the opponent's choice. What if uncertainty arises from, for example, increased perceptual ambiguity? Further research is needed to determine whether the rIPL is also involved in resolving perceptual ambiguity.

## Funding

Office of Naval Research YIP2005 award (N000140810407) to Y.V.J.; and National Science Foundation graduate fellowship to T.J.V.

## Notes

*Conflict of Interest:* None declared.

Address correspondence to Timothy J. Vickery, Department of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520, USA. Email: tim.vickery@gmail.com.

## References

- Attneave F. 1959. Applications of information theory to psychology. New York: Henry Holt and Company.
- Baddeley AD. 1966. The capacity for generating information by randomization. *Q J Exp Psychol.* 18:119-129.
- Baddeley A, Emslie H, Kolodny J, Duncan J. 1998. Random generation and the executive control of working memory. *Q J Exp Psychol A.* 51:819-852.
- Barracough DJ, Conroy ML, Lee D. 2004. Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci.* 7:404-410.
- Bechara A, Damasio H, Tranel D, Anderson SW. 1998. Dissociation of working memory from decision making within the human prefrontal cortex. *J Neurosci.* 18:428-437.
- Brainard DH. 1997. The psychophysics toolbox. *Spat Vis.* 10:433-436.
- Camerer CF. 2003. Behavioral game theory: experiments on strategic interaction. Princeton (NJ): Princeton University Press.
- Carter CS, MacDonald AM, Botvinick MM, Stenger VA, Cohen JD. 2000. Parsing executive processes: control vs. evaluation in the frontal cortex. *Biol Psychiatry.* 47:104S-104S.
- Chafee MV, Averbeck BB, Crowe DA. 2007. Representing spatial relationships in posterior parietal cortex: single neurons code object-referenced position. *Cereb Cortex.* 17:2914-2932.
- Dale AM. 1999. Optimal experimental design for event-related fMRI. *Hum Brain Mapp.* 8:109-114.
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. 2006. Cortical substrates for exploratory decisions in humans. *Nature.* 441:876-879.
- Dorris MC, Glimcher PW. 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron.* 44:365-378.
- Duncan J, Owen AM. 2000. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* 23:475-483.
- Fel'tovich N. 2000. Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica.* 68:605-641.

- Fudenberg D, Levine DK. 1998. *The theory of learning in games*. Cambridge (MA): MIT Press.
- Glimcher PW. 2005. Indeterminacy in brain and behavior. *Annu Rev Psychol*. 56:25-56.
- Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG. 2004. A general mechanism for perceptual decision-making in the human brain. *Nature*. 431:859-862.
- Huettel SA, Song AW, McCarthy G. 2005. Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *J Neurosci*. 25:3304-3311.
- Kanwisher N, Wojciulik E. 2000. Visual attention: insights from brain imaging. *Nat Rev Neurosci*. 1:91-100.
- Lee D, Conroy ML, McGreevy BP, Barraclough DJ. 2004. Reinforcement learning and decision making in monkeys during a competitive game. *Brain Res Cogn Brain Res*. 22:45-58.
- Maunsell JHR. 2004. Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci*. 8:261-265.
- Miller MB, Valsangkar-Smyth M, Newman S, Dumont H, Wolford G. 2005. Brain activations associated with probability matching. *Neuropsychologia*. 43:1598-1608.
- Neuringer A. 1986. Can people behave randomly—the role of feedback. *J Exp Psychol Gen*. 115:62-75.
- Nickerson RS. 2002. The production and perception of randomness. *Psychol Rev*. 109:330-357.
- Paulus MP, Hozack N, Zauscher B, McDowell JE, Frank L, Brown GG, Braff DL. 2001. Prefrontal, parietal, and temporal cortex networks underlie decision-making in the presence of uncertainty. *Neuroimage*. 13:91-100.
- Pelli DG. 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis*. 10:437-442.
- Rapoport A, Budescu DV. 1992. Generation of Random series in 2-person strictly competitive games. *J Exp Psychol Gen*. 121:352-363.
- Rushworth MFS, Walton ME, Kennerley SW, Bannerman DM. 2004. Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci*. 8:410-417.
- Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol*. 86:1916-1936.
- Shuman M, Kanwisher N. 2004. Numerical magnitude in the human parietal lobe: tests of representational generality and domain specificity. *Neuron*. 44:557-569.
- Sugrue LP, Corrado GS, Newsome WT. 2004. Matching behavior and the representation of value in the parietal cortex. *Science*. 304:1782-1787.
- Sutton RS, Barto AG. 1998. *Reinforcement learning: an introduction*. Cambridge (MA): MIT Press.
- Towse JN, Valentine JD. 1997. Random generation of numbers: a search for underlying processes. *Eur J Cogn Psychol*. 9:381-400.
- Walton ME, Bannerman DM, Alterescu K, Rushworth MFS. 2003. Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *J Neurosci*. 23:6475-6479.