



ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage xx (2007) xxx–xxx

Nonlinear neurobiological probability weighting functions for aversive outcomes

Gregory S. Berns,^{a,*} C. Monica Capra,^b Jonathan Chappelow,^a
Sara Moore,^a and Charles Noussair^b

^aDepartment of Psychiatry and Behavioral Sciences, Emory University School of Medicine, 101 Woodruff Circle, Suite 4000, Atlanta, GA 30322, USA

^bDepartment of Economics, Emory University, Atlanta, GA 30322, USA

Received 1 June 2007; revised 15 October 2007; accepted 21 October 2007

While mainstream economic models assume that individuals treat probabilities objectively, many people tend to overestimate the likelihood of improbable events and underestimate the likelihood of probable events. However, a biological account for why probabilities would be treated this way does not yet exist. While undergoing fMRI, we presented individuals with a series of lotteries, defined by the voltage of an impending cutaneous electric shock and the probability with which the shock would be received. During the prospect phase, neural activity that tracked the probability of the expected outcome was observed in a circumscribed network of brain regions that included the anterior cingulate, visual, parietal, and temporal cortices. Most of these regions displayed responses to probabilities consistent with nonlinear probability weighting. The neural responses to passive lotteries predicted 79% of subsequent decisions when individuals were offered choices between different lotteries, and exceeded that predicted by behavior alone near the indifference point.

© 2007 Elsevier Inc. All rights reserved.

Introduction

Most decisions that people make involve some element of uncertainty. When the decision is between alternatives that result in different distributions of potential outcomes, the individual faced with the decision must consider the potential benefit of each outcome, as well as the likelihood of its occurrence. In traditional economic analysis, individuals are assumed to make decisions to maximize their *expected utility*. The expected utility of a lottery is the benefit (or utility) of each possible outcome multiplied by the probability of its occurrence (von Neumann and Morgenstern, 1944). Although this assumption leads to a theoretically parsimonious framework for analyzing decision making, behavioral experiments have accumulated considerable evidence against the

hypothesis that individuals make decisions as if they weight the utility of each outcome by its objective probability. These results have, in part, led to the formulation of alternatives to expected utility theory, such as prospect theory, rank-dependent utility theory, and cumulative prospect theory, in which probabilities are treated in a nonlinear manner (Allais, 1953; Kahneman and Tversky, 1979; Starmer, 2000). Although probability weighting explains a wide range of empirical findings, it is unknown why people behave in this manner, and the underlying neurological mechanisms at work are not understood. In this paper, we provide experimental evidence and a biological rationale for this pattern of probability weighting.

Utility functions were originally proposed as a theoretical construct to mathematically summarize preferences over outcomes. This facilitated the solution of economic models employing the assumption that economic actors made decisions with the purpose of attaining preferred outcomes. Neuroeconomists, in contrast, have recently advanced the hypothesis that utility is a neural response and that activations of certain brain regions provide measures of utility, in the sense that individuals make decisions as if they are maximizing the activation of these regions (Camerer et al., 2005; Glimcher, 2002; Hsu et al., 2005; Huettel et al., 2006; Knutson et al., 2005; Preuschoff et al., 2006). We take a similar approach here. We measure how a stimulus, in the form of a lottery consisting of aversive outcomes and associated probabilities, is transformed into a neurobiological response. We then propose a neurobiological measure of probability weighting, called the neurobiological probability response ratio (NPRR). According to this measure, the neurobiological responses of participants to the presentation of risky lotteries are consistent with nonlinear probability weighting. We then consider whether observed levels of neurobiological responses to different lotteries provide accurate predictions of individuals' choices between the lotteries. We obtain evidence that individuals act to minimize the neurobiological measures associated with losses and that the heterogeneity in neurobiological measures among subjects explains much of the individual differences in decisions.

* Corresponding author. Fax: +1 404 727 3233.

E-mail address: gberns@emory.edu (G.S. Berns).

Available online on ScienceDirect (www.sciencedirect.com).

We performed functional magnetic resonance imaging (fMRI) and skin conductance response (SCR) measurements on 37 participants while they were presented with a series of lotteries (Fig. 1). Each lottery, or prospect, defined the probability and magnitude of an imminent electrical shock to the left foot. The magnitude of the shock ranged from 10% to 90% of each individual's stated maximum voltage (where 0% is the minimum voltage perceptible to the subject), and the probabilities ranged from 1/6 to 1. In the first (passive) phase of the experiment, participants did nothing other than observe the prospect, wait 8 s, and then experience the outcome (shock or no shock). They subsequently rated the experience using a visual analogue scale (VAS). Because no decision was required, the fMRI activations obtained in this phase were representative of intrinsic responses to voltage and probability and, by inference, the expected utility of the lottery. We study the degree to which the implied neurobiological probability weighting functions are consistent with non-linear probability weighting. In the second (active) phase, participants were offered choices between two lotteries. We then considered the extent to which decisions taken in the active phase

are consistent with attempting to minimize the values of the fMRI activation observed in the passive phase.

Methods

Subjects

A total of 37 (20 females, 17 males) people were scanned using fMRI. We observed a higher than usual rate of signal artifacts that necessitated the discarding of 9 subjects, leaving $N=28$ for the data analysis (14 males, 14 females; ages: 19–42). In part, this was due to scanner gradient malfunctions, but the complexity of the setup with shock electrodes and SCR leads also contributed to the high rate of artifacts. The decision to discard was based on the appearance of more than 5 spikes in the mean intensity of the fMRI signal for the whole brain. Skin conductance responses (SCRs) were collected on all participants (although only 26 of these were usable due to signal artifacts from the scanner). Each participant gave written, informed consent for a protocol approved by Emory University's IRB. Each participant's session lasted for an

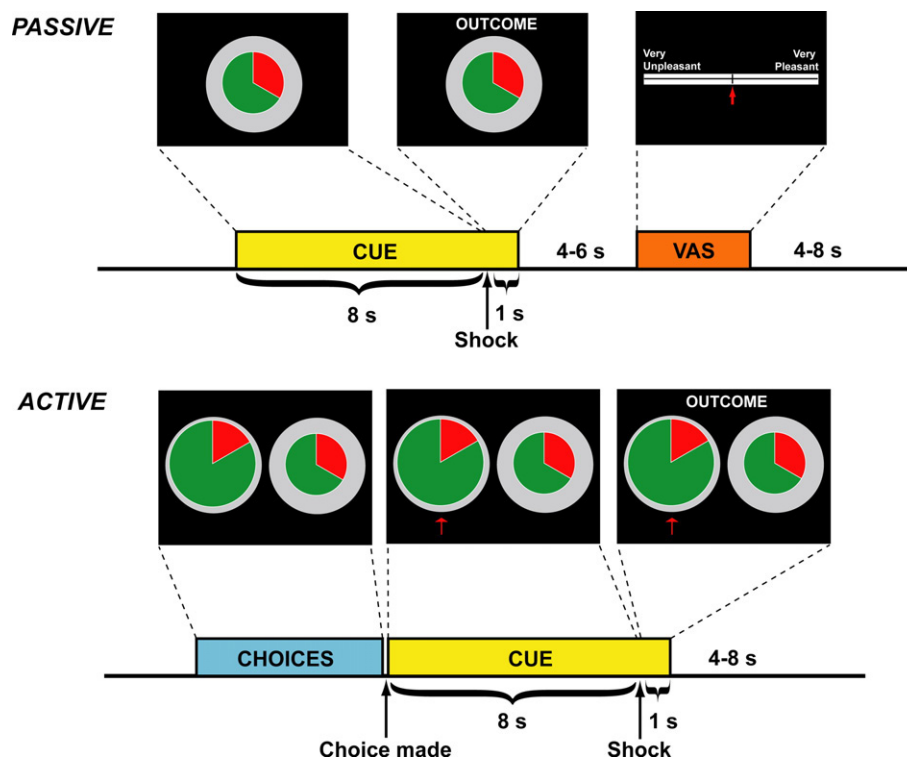


Fig. 1. Functional MRI trial design. Each trial followed a delay-conditioning design (*PASSIVE*), in which a cue was presented for the duration of the trial, up to, and beyond, the delivery of an aversive stimulus in the form of a brief cutaneous electric shock (10–15 ms duration). At the beginning of each trial, a cue was displayed that indicated the magnitude of the shock and the probability with which that shock would be delivered. The magnitude was indicated by the size of the colored circle relative to the outer gray circle, which denoted the individual's maximum tolerable voltage. Four voltage levels were used [10%, 30%, 60% (*shown*), and 90% of the maximum]. The probability was indicated by the proportion of the inner circle colored *red*. Five probabilities were used [1/6, 1/3 (*shown*), 2/3, 5/6, and 1]. All 20 combinations of voltage and probability were administered in a repeated event design. The word, "OUTCOME," was presented simultaneously with the delivery of the shock on those trials in which a shock occurred, and it also appeared at the same point in time on those trials in which a shock was not delivered, providing an indicator that the trial was over and the participant would not receive a shock in the current trial. Following the shock, the cue remained visible for another 1 s to prevent conditioning to the cue offset. A visual analog scale (VAS) was then presented in which the individual moved an arrow to evaluate her subjective experience for the entire preceding trial, including the waiting time. Following three passive runs, consisting of 40 trials each, a choice run was presented (*ACTIVE*). Here, each trial consisted of two voltage–probability gambles, and the participant had to choose between them. Following each choice, the outcome was determined by a random number chosen from the appropriately weighted binomial distribution. The choices offered all had the property that one alternative corresponded to a higher voltage, as well as a lower probability; e.g. 1/6 chance of 90% of maximum shock vs. 1/3 chance of 60% of maximum (*shown*). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

average of 2 h, and each individual was awarded \$40 at the end of his session.

Experimental procedures

Cutaneous electrical shocks were delivered using a Grass SD-9 stimulator (West Warwick, RI) through shielded, gold electrodes placed 2–4 cm apart on the dorsum of the left foot. Each shock was a monophasic pulse of 10- to 20-ms duration. The Grass stimulator was modified by attaching a servo-controlled motor to the voltage potentiometer. The motor allowed for computer-control of the voltage level without compromising the safety of the electrical isolation in the stimulator. The motor was controlled by a laptop through a serial interface.

Prior to scanning, the voltage range was titrated for each participant. The detection threshold was determined by delivering pulses starting at zero volts and increasing the voltage until the individual indicated that he could feel them. This minimum perception threshold is denoted as V_{\min} in this paper for purposes of exposition. The voltage was increased further, while each participant was instructed, “When you feel that you absolutely cannot bear any stronger shock, let us know – this will be set as your maximum; we will not use this value for the experiment, but rather to establish a scale. You will never receive a shock of maximum value”. This maximum is denoted as V_{\max} . The purpose of this procedure was to control for the heterogeneity of the skin resistance among subjects and to administer potentially painful stimuli in an ethical manner. We measure the strength of the shock administered to an individual by s , where the associated voltage for an individual is $V_s = s \times V_{\max} + (1-s) \times V_{\min}$. For the remainder of the experiment, s took on values of 10%, 30%, 60%, and 90%.

After the voltage titration, the scanning phase began. The software package, COGENT 2000 (FIL, University College London), was used for stimulus presentation and response acquisition. The first phase of the experiment, which we call the *Passive Phase*, consisted of 120 trials. At the beginning of each trial, each participant was presented with a pie chart that conveyed both the voltage of the impending shock and the probability with which it would be received (Fig. 1, *Passive*). The display thus defined a lottery consisting of a shock with voltage level, s , and a probability, p , in trial, t (s^t, p^t). The size of the pie chart indicated the strength of the shock that might be applied in the current trial, with the area of the pie chart equaling s times the area of the outer reference circle (denoting V_{\max}). The percentage of the inner circle that was filled in red indicated the probability with which the shock was to be administered. The possible probabilities were 1/6, 1/3, 2/3, 5/6, and 1. With the four voltage levels, this yielded 20 voltage–probability combinations, each of which was presented 6 times during the 120 trials that made up the passive phase of a session.

In each of the 120 trials of the passive phase, the individual observed the display indicating the voltage–probability parameters in effect for the trial. After 8 s, the shock was then applied with the appropriate probability. The order of the outcomes was predetermined to ensure that there would be at least one event in each of the 20 conditions (four voltage levels times five probability levels) in each of the three scan runs. Although the order was predetermined, the frequency of shocks received in each of the conditions reflected the actual probability. There were 6 repetitions of each of the 20 voltage–probability combinations, and so, for example, under the lowest probability, 1/6, a shock was administered only once at each of the voltages. After the realization

of the outcome, whether or not a shock actually occurred, the subject was required to rate the experience of the trial, in a range between “very unpleasant” and “very pleasant.” To indicate his rating, he marked a location on a visual analog scale (VAS), using a cursor operated with his hand control (Noussair et al., 2004).

In the second phase, which we call the *Active Phase*, each individual faced a sequence of 60-pairwise choices from the set of probability/shock combinations presented in the passive phase. In each round, two displays similar to the one shown in Fig. 1 (*Active*) appeared side by side, and subjects were required to choose one of the two lotteries, using the keypad provided to them. The experimenter chose the pairs so that one alternative always specified both a higher voltage shock as well as lower probability than the other alternative. In other words, if s^t_A and s^t_B are the strengths of the shocks that may be applied under choices A and B , respectively, in trial t , it was always the case that $s^t_A > s^t_B \leftrightarrow p^t_A < p^t_B$, where p^t_A is the probability that a shock is applied under choice A in trial t . The 60 choices represented all of the possible combinations meeting this voltage–probability constraint. Subjects received a shock with the voltage specified by their choice and with the indicated probability. Outcomes were determined on-the-fly by a random number drawn from the appropriate binomial distribution.

fMRI measurements

Scanning was conducted with a Siemens 3 T Trio whole-body scanner. After acquisition of a high-resolution T1-weighted scan, fMRI of the BOLD response was performed (TR=2350 ms, TE=30 ms, 64×64 matrix, 35 axial slices, 3 mm^3 cubic voxels). To prevent electrical artifacts in the fMRI signal due to shock deliveries, the latter occurred during a 50-ms pause after each volume, yielding an effective TR=2400 ms. Three runs of 40 trials were performed during the passive phase, and one run of 60 trials for the active phase, for a total scan time of approximately 75 min.

SCR measurements

Skin conductance responses (SCR) were acquired simultaneously with the scanning with a Biopac MP150 digital converter (Biopac Systems, Goleta, CA) and fed into AcqKnowledge 3.7 recording software. A TTL-generator box was interfaced to the serial port of the computer running COGENT, allowing for the generation of digital timestamps for each stimulus on the Biopac channel recordings. The SCR data were sampled at 125 Hz, and a 1 Hz low-pass filter and 0.05 Hz high-pass filter were applied to the data during acquisition. The SCR data were first detrended and spike artifacts from the scanner removed. Average responses were computed in a peristimulus window aligned to the shock. We used the integrated amplitude from the cue onset to the shock as a measure of the cue-related SCR and the integrated amplitude from 2 to 10 s after the shock as a measure of the shock-related SCR.

fMRI analyses

fMRI data were analyzed with SPM2 with random-effects models. Standard preprocessing was used, including motion correction, slice timing correction, and normalization to the MNI template brain in Talairach orientation. In the passive phase, each trial was modeled as a 9-s variable duration cue, and the shock was modeled as an impulse function. Similarly, the outcomes in which no shock was delivered were also modeled as impulse functions.

Each of the 20 combinations of voltage and probability was considered as different, and thus there were 60 conditions. The 60 conditions consisted of 20 probability–shock combinations, times three states: presentation of cue, realized shock, and realized avoidance of a shock. Regressors for the VAS and motion parameters were also included.

To identify brain regions that were associated with either probability or voltage magnitudes, we used two simple linear contrasts. The voltage weighting contrast identified regions that responded in a monotonically increasing fashion to the prospect of higher voltage outcomes. To isolate the effect of magnitude, we specified a contrast using only those trials with a certain outcome (probability 1 of a shock). This contrast used a linear weighting function for voltage (0.1, 0.3, 0.6, 0.9) centered around the mean. In a similar fashion, to identify probability sensitive regions, we used a linear contrast in probability (1/6, 2/6, 4/6, 5/6, 1; centered around the mean) for the highest voltage prospect. We restricted the probability contrast to the highest voltage prospect because this was the most salient outcome and therefore the most likely to provide an accurate map of brain regions responsive to probability. Both maps were thresholded at $P < 0.001$ (uncorrected) and $k > 10$. ROIs (6 mm radius spheres) were drawn on these functional maps and the average beta values for the cue in the 20 voltage–delay combinations were extracted for each participant in these ROIs. Subsequent analyses were performed outside SPM.

Probability weighting functions

In order to determine a probability weighting function with fMRI measurements, we must define a metric that has certain properties. The probability weighting function, $w(p)$, is a monotonic function that transforms an objective probability, p , into a “weight” that is applied when evaluating a lottery. It is bounded by $[0, 1]$ in both domain and range such that $w(0) = 0$ and $w(1) = 1$. Let $y(x, p)$ represent the fMRI response of a particular brain region to a lottery with outcome x and probability p (and a null or status quo outcome with probability $1 - p$). Then the neurobiological equivalent of a probability weighting function is given by the ratio: $y(x, p)/y(x, 1)$, where $y(x, 1)$ is the response to a lottery with outcome x and probability 1. We refer to this ratio as the *neurobiological probability response ratio* (NPRR). We hypothesize that the form of this ratio is directly related to the form of probability weighting observed in previous empirical studies and may provide a neurobiological basis for probability weighting. This definition of the NPRR has the property of being separable in utility and probability.

However, many biological measurements, such as those obtained with fMRI, do not represent absolute physical quantities, and as such, must be interpreted as relative measurements. Thus, to calculate the NPRR, we must define an effective baseline, y_0 , from which to reference neurobiological responses, and specify the NPRR as:

$$\text{NPRR}(x, p) = \frac{[y(x, p) - y_0]}{[y(x, 1) - y_0]} \quad (1)$$

Both the value function for shock prospects and the probability weighting function were determined in the ROIs that activated in response to the stimulus, the presentation of the lotteries. The value function, $y(x, 1)$, was determined by the mean beta value of

the BOLD activation as a function of voltage for the trials in which the probability of a shock was 1. This was the condition of no uncertainty and was the most appropriate for the determination of the relationship between actual and perceived shock voltages. The choice of y_0 , however, was more problematic. In order to control for various psychological and physiological processes associated with low-level processing of the stimuli themselves, y_0 should be as similar as possible to the lottery conditions, which rules out using an implicit baseline (null) condition. The choice of y_0 will affect the shape of the NPRR relative to the diagonal but not the relative orderings of NPRR. An extensive behavioral literature on probability weighting in the financial domain suggests that the probability weighting function crosses the diagonal in the vicinity of $p = 0.4$ (Abdellaoui, 2000; Fehr-Duda et al., 2006; Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Wu and Gonzalez, 1996). In our experiment, the probability prospect that came closest to meeting these requirements was $p = 1/3$. Thus, probability weighting functions were calculated according to Eq. (1), where y_0 was the response to the probability prospect ($p = 1/3$) at a given voltage. For direct comparison to linear probability weighting, we referenced the $p = 1/3$ condition to its objective probability. This resulted in the following functional form for the NPRR, which was used for the data analysis in this study:

$$\text{NPRR}(x, p) = \frac{2[y(x, p) - y(x, 1/3)]}{3[y(x, 1) - y(x, 1/3)]} + \frac{1}{3} \quad (2)$$

Because of deviations from normality in the distributions of NPRR, namely a few large-valued outliers resulting from the division operation, we used the median of the means of the three highest voltages value at each probability to estimate the central tendency. 95% confidence intervals of the median were computed according to this formula: $[(N+1)/2] \pm 1.96 \times (\sqrt{N})/2$, which corresponded to subjects 9 and 20 in an ordered listing of the 28 participants (Woodruff, 1952). To directly test the null hypothesis of linear probability weighting, e.g. EUT, we used the 95% confidence intervals at $p = 1/6, 2/3, \text{ and } 5/6$. Because previous behavioral evidence suggests nonlinear probability weighting is most prominent at small probabilities, we were particularly interested in whether $\text{NPRR}(1/6) > 1/6$, which would reject the null hypothesis in favor of overweighting of small probabilities. The same analysis was performed on the SCR data, except that y_0 was taken to be 0 because the SCR data were already referenced to zero by virtue of the temporal filtering and integration procedure.

Linking passive activation to decision making

To determine whether the probability weighting functions obtained from the passive trials could predict decision making in the active phase, we performed a logistic regression on the participants' choices during the active trials. Each active trial contained two lotteries, designated *left* (L) and *right* (R), based on their location on the display shown to participants. We hypothesized that an individual would choose the lottery to minimize some objective function $V(x, p)$. If we take the passive neurobiological activations in the ROIs as a proxy for $V(x, p)$, then for each pair of lotteries over which an individual must choose, there are two expected values for each ROI: $y(x_L, p_L)$ and $y(x_R, p_R)$, for the *left* and *right* lotteries, respectively. We hypothesized that, for a given

choice pair, the probability of choosing *right* is given by the logistic relationship:

$$\text{logit}(P_R) = \ln\left(\frac{P_R}{1-P_R}\right) = \beta_0 + \sum_i \beta_i [y_i(x_{R,P_R}) - y_i(x_{L,P_L})] \quad (3)$$

where P_R is the probability of choosing the right-hand lottery and i is the i th ROI. Negative coefficients are consistent with the individual choosing the lottery with the lower passive activation.

In order to determine whether Eq. (3) is a good predictor of choice, we performed a bootstrap logistic regression on the choice data (Davison and Hinkley, 1997). The bootstrap procedure performed the logistic regression 10,000 times for each individual. Each participant made choices over 60 pairs of lotteries, and in each iteration of the bootstrap, we randomly chose, without replacement, 40 of these trials. These 40 trials were then used to determine the coefficients, β_i , in Eq. (3). As inputs to Eq. (3), we used the mean activation in each ROI obtained in the passive procedure for the two lotteries that corresponded to the left and right lotteries offered in the choice pair. For example, if the left-hand lottery was 60% voltage with 1/3 probability, and the right-hand lottery was 30% voltage with 5/6 probability, we took the mean passive activation for these two lotteries and computed the difference for each ROI. This vector of activation differences across the ROIs comprised the inputs on the right side of Eq. (3). We used a subset of the ROIs listed in Table 1. The activations of some of the ROIs were correlated with each other, which could result in an overspecification for the logit regression. Three pairs of ROIs were highly correlated (L and R insula/superior temporal, R superior frontal and L cingulate, precuneus and L postcentral), and so the mean activation for these pairs in each condition and each subject were used as inputs to Eq. (3). Coefficients were fit for each subject separately. The 20 trials not used in each iteration served as test trials for the accuracy of the corresponding regression. On a given test trial, the model error was measured as the absolute difference between the model output, P_R , and the actual choice (*left*=0, *right*=1). Thus, the accuracy was given by: $1 - |P_R - \text{choice}|$.

In order to compare these results to those of a model employing only decisions and not brain activations, we computed a rank ordering for all of the lotteries based on the choices observed for each 40 trials in each bootstrap iteration. Because only a subset of all possible pairings was presented, we assumed that for identical probabilities, an individual would have always chosen the lottery with a lower expected voltage. Similarly, we assumed that for lotteries with identical voltages, he would have always chosen the lottery with a lower probability of shock. Finally, we assumed that a lower voltage/lower probability lottery would have been chosen over a higher voltage/higher probability lottery. These assumptions, plus the actual choices on trials that traded off voltage for probability, allowed us to calculate a rank ordering of all 20 lotteries. The ranks were computed using the Colley matrix method (Berns et al., 2006; Colley, 2002). In order to compare the ability of this behaviorally derived ranking to the passive brain activation, we then used these ranks as inputs into a logistic regression:

$$\text{logit}(P_R) = \ln\left(\frac{P_R}{1-P_R}\right) = \beta_0 + \beta [R(x_{R,P_R}) - R(x_{L,P_L})] \quad (4)$$

Table 1

Brain regions showing a correlation of activity during passive phase to lottery cue signaling outcome magnitude (cluster size >10 with $P < 0.001$)

Region	MNI coordinates	T statistic	NPRR(1/6) > 1/6 ?
<i>Magnitude-sensitive regions (positive correlation)</i>			
Visual cortex (BA 17)	9, -84, 3	6.70	Y
R Superior frontal gyrus (BA 6)	6, 3, 63	6.06	N
Bilateral cerebellum	-27, -66, -21	5.94	Y
L Insula/superior temporal gyrus (BA 38)	-54, 15, -9	5.15	Y
R Precentral gyrus (BA 6)	48, -6, 60	5.15	Y
R Insula/superior temporal gyrus (BA22)	60, 12, -3	5.06	Y
Precuneus (BA 7)	0, -54, 57	4.90	Y
<i>Magnitude-sensitive regions (negative correlation)</i>			
R Visual cortex (BA 18)	24, -93, 0	6.40	N
L Caudate	-30, -45, 12	6.02	Y
L Precentral gyrus (BA 6)	-30, 0, 36	5.93	Y
L Cuneus (BA 17)	-18, -99, -6	5.36	Y
R Cingulate gyrus (BA 24)	21, 3, 39	5.29	Y
R Parietal	27, -36, 42	4.96	Y
R Middle frontal gyrus (BA 11)	33, 45, -6	4.91	Y
L Thalamus	-15, -33, 6	4.87	N
L Inferior parietal lobule (BA 40)	-45, -39, 48	4.81	Y
L Posterior cingulate (BA 31)	-18, -45, 36	4.45	N
R Posterior insula (BA 13)	36, -21, 24	4.44	Y
R Postcentral gyrus (BA 40)	48, -36, 51	4.35	Y
R Precuneus (BA 31)	9, -60, 30	4.30	Y
L Inferior temporal gyrus (BA 21)	-63, -9, -21	4.15	Y
R Caudate/ACC	9, 21, -3	4.07	Y
R Cingulate gyrus (BA 32)	21, 21, 30	3.88	N
R Inferior parietal lobule (BA 39)	48, -69, 42	3.78	N

Regions showing significant ($P < 0.05$) departures from linear probability weighting at probability=1/6 are indicated in the rightmost column (none were significant for $\text{NPRR}(1/6) < 1/6$).

where $R(x_{L,P_L})$ and $R(x_{R,P_R})$ were the ranks of the left- and right-hand lotteries. Like the passive fMRI model, the accuracy was measured as $1 - |P_R - \text{choice}|$.

Results

From the selections obtained during the choice procedure, we calculated a relative ranking for the 20 combinations of voltage and probability, and thereby obtained a preference ordering over all 20 lotteries. The rankings represented a linear transformation of the percentage of time a given lottery was chosen from all of the instances in which it was offered. Both voltage [$F(3,144)=404$, $P < 0.0001$] and probability [$F(4,144)=406$, $P < 0.0001$] significantly affected a given lottery's rank, with higher voltage and higher probability lotteries being ranked worse.

We used several approaches to test whether participants exhibited behavioral evidence of nonlinear probability weighting (Berns et al., 2007). First we used the lottery ranks to estimate

values for γ in a standard specification for nonlinear weighting (Tversky and Kahneman, 1992):

$$V(L) = \sum_i -\lambda \frac{P_i^\gamma}{(\sum_i P_i^\gamma)^{1/\gamma}} |X_i|^\alpha \quad (5)$$

which, across all 37 subjects, yielded a median value of 0.685 for γ and was consistent with an inverted S-shaped probability weighting function. We also considered the incidence of “common-ratio-violations.” In our experiment, common ratio violations were observed when the lottery ($V_h, p_h=1/6$) was chosen over ($V_l, p_l=2/6$), but ($V_l, p_l=4/6$) was chosen over ($V_h, p_h=2/6$), or vice versa. Out of six possible instances, the average number of common ratio violations was 1.95 (SEM 0.23), which was significantly different from zero and not generated by random choice (the Friedman test; $\chi^2=10.80$; $df=1$; $p<0.001$). Under the assumption of linear probability weighting, multiplying the probability of both lotteries by a constant should not affect choices; that is, we should not expect any violations. On the other hand, common ratio violations could be a consequence of probability weighting. In particular, an individual whose $w(p)$ has the inverted S-shape, would exhibit more risk aversion over small probabilities of larger losses, resulting in indifference curves in the probability triangle that fan out. In our experiments, we

observed that, in general, the direction of the violations were consistent with fanning out of indifference curves (Berns et al., 2007).

In the passive phase, a well-defined network of brain regions was active in response to the cue, and the pattern of activation could be largely dissociated into magnitude-sensitive and probability-sensitive regions (Fig. 2 and Tables 1 and 2). Some of these regions were clearly related to the low-level processing of visual stimuli (e.g. visual cortex). Other regions, however, encoded aspects of the anticipated voltage and/or probability. There was very little anatomical overlap between the magnitude and probability maps, with the exception of the ACC and supplementary motor area dorsal to it. The probability of receiving a shock was most significantly correlated with activity in the bilateral inferior parietal cortex, near the temporal–parietal–occipital junction (Fig. 2, green), whereas the magnitude of the impending shock was correlated with bilateral activity in the insula/superior temporal cortex, precuneus, cerebellum, and a region of the precentral gyrus associated with the foot (Fig. 2, red). Many other regions, including the caudate/subgenual ACC, displayed negative correlations with the impending magnitude of the outcome or the probability of receiving a shock. When compared to the implicit baseline, however, all of these regions were deactivated. Interestingly, these too, displayed nonlinear forms of probability weighting (see below).




Fig. 2. Brain regions responsive to cue indicating either magnitude (red) or probability (green) of an impending electrical shock in the passive phase ($P<0.001$, cluster size >10). Magnitude sensitive regions were determined by a linearly weighted contrast by voltage of certain prospects ($p=1$). Probability sensitive regions were identified by a linearly weighted contrast by probability of the highest voltage prospects ($V=90\%$). The networks for expected magnitude and probability were largely separable with occipital and parietal regions being responsive to magnitude. Probability-specific regions included bilateral parietal–occipital junction and postcentral gyrus for the foot. The only region that was responsive to both expected magnitude and probability was an area between the ACC and SMA (yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
Brain regions showing a correlation of activity during passive phase to lottery cue signaling outcome probability (cluster size > 10 with $P < 0.001$)

Region	MNI coordinates	T statistic	NPRR(1/6) > 1/6 ?
<i>Probability-sensitive regions (positive correlation)</i>			
R Cingulate gyrus (BA 24)	9, -18, 39	5.88	N
L Inferior parietal gyrus (BA 40)	-57, -57, 42	4.98	N
R Superior temporal gyrus (BA 22)	63, -48, 12	4.97	Y
R Superior frontal gyrus (BA 6)	9, 3, 63	4.92	Y
R Inferior parietal gyrus (BA 40)	66, -36, 30	4.91	N
L Cingulate gyrus (BA 24)	-6, 0, 45	4.89	N
L Postcentral gyrus (BA 7)	-9, -57, 69	4.64	Y
<i>Probability-sensitive regions (negative correlation)</i>			
L Visual cortex (BA 18)	-30, -96, 6	5.58	N
L Posterior cingulate (BA 29)	-12, -51, 6	4.66	Y
R Lingual gyrus (BA 19)	18, -57, -3	4.30	N
L Anterior cingulate (BA 24)	-6, 21, -3	4.05	Y
R Thalamus	21, -30, -3	4.03	Y
L Middle temporal gyrus (BA 37)	-48, -60, -3	4.00	Y

Regions showing significant ($P < 0.05$) departures from linear probability weighting at probability = 1/6 are indicated in the rightmost column (none were significant for $\text{NPRR}(1/6) < 1/6$).

Within these regions, the relationship of activation to probability defines a neurobiological probability response ratio (NPRR). The NPRR can be viewed as a neurobiological analogue to the probability weighting function, $w(p)$, where p is the objective probability of an outcome and $w(p)$ is the transformed probability that is consistent with an individual's decisions. Prospect theory assumes that $w(p)$ has an inverted S-shape, overweighting moderately low probabilities and underweighting moderately high ones, while expected utility theory assumes that $w(p) = p$. While prospect and expected utility theories were proposed as models of decision behavior, we take the view that any theory that describes decisions must have an underlying biological foundation. Thus, *we will interpret the linearity of the NPRR functions as evidence in favor of expected utility theory, and consider inverted S-shaped NPRR functions as evidence for prospect theory.* As shown in Fig. 3, most regions displayed a noticeable nonlinearity. An inverted S shape form, as postulated in prospect theory, was the most common. To statistically test this nonlinearity, the value of NPRR (1/6) was compared to 1/6 (i.e. linear weighting). If NPRR(1/6) was significantly greater than 1/6, this was taken as support for nonlinear weighting. As shown in Tables 1 and 2, the majority of regions were significantly nonlinear (25 of the 37 regions listed, the null hypothesis that this was a random property could be rejected with $\chi^2 = 4.6$, $df = 1$, $P = 0.03$). The integrated SCRs, however, were not significantly different from a linear relationship to probability (Fig. 4).

The overall pattern of brain activity obtained during the passive phase suggested that the intrinsic response to risky prospects was generally nonlinear in probability. In order to determine whether these passive brain responses are consistent with later decision making, we performed a logistic regression on the participants' choices. Using a bootstrap procedure in which 2/3 of the choice trials were used for the regression and 1/3 for testing its accuracy, we found that the difference in passive activation for the two lotteries accurately predicted individuals' choices for 79% (versus an overall accuracy of 83% for behavioral data alone, but this was

not significantly different on a subject-wise paired t -test, $P = 0.11$) of the test trials. The prediction accuracy of the brain activation approached 100% when the two offerings were sufficiently far apart in rank (Fig. 5).

Interestingly, the predictive power of passive brain activation exceeded the predictive power of individual's own average choice behavior, as reflected in the rankings calculated from his decisions, when the individual was close to indifferent between the available alternatives. Because the lottery ranks were determined from the entire set of choices, the ranks represented an average ordering, but because some individuals were inconsistent in their choices, these rankings were not able to fully specify choices on individual trials (except when the choices were sufficiently different in rank). When the ranking difference between lotteries was less than 3, the passive brain activation was significantly better at predicting choice than the rankings themselves ($P < 0.01$).

With regard to ex-post probability and voltage strength, measures of the subjective response yielded general patterns that were consistent with the brain imaging and SCR results. On trials in which a shock was received, the post-outcome VAS self-reported ratings of the experience of the trial indicated a significantly positive effect of both voltage [$F(3,144) = 131$, $P < 0.0001$] and probability [$F(4,144) = 4.1$, $P = 0.003$]. However, on trials in which no shock was received, neither voltage [$F(3,108) = 0.51$, $P = 0.68$] nor probability [$F(3,108) = 0.92$, $P = 0.43$] had a significant effect on the VAS rating. Thus, while the experienced disutility of the shock, as measured in participants' self-reports, was dependent on the voltage and the probability (with higher probability trials being rated as slightly worse), the utility of avoiding a shock was not.

Discussion

The way that probability enters into decision making has been an ongoing area of controversy in economics (Starmer, 2000). Expected utility theory (EUT) assumes that probabilities are treated linearly and that decisions are based on the utility of each possible outcome, weighted by the probability of its occurrence (Savage, 1954; von Neumann and Morgenstern, 1944). Although EUT remains a foundational assumption in most fields of economics, as well as in game theory, its descriptive accuracy for actual behavior has been repeatedly questioned (Allais, 1953; Kahneman and Tversky, 1979; Wu and Gonzalez, 1996). Several alternatives, including the widely accepted prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992), suggest that observed patterns of decision making under risk can, in part, be explained if individuals use probabilities in a nonlinear manner. This principle seems to be true for both gains and losses. However, a biological explanation for why this would occur does not yet exist. Here, we suggest that the source of nonlinear probability weighting lies in how the brain transforms representations of probability into behaviorally relevant heuristics.

When probability information is conveyed visually, whether transmitted as a number or as a graphical representation, the information must be transformed into a mental construct that represents the likelihood of an outcome. Although this transformation can be thought of as a psychological process, whatever mental processes are involved are constrained by the biological properties of the brain that implement the processes. Our SCR data demonstrate a monotonic relationship to the probability of receiving a shock, which is consistent with a rich literature linking autonomic arousal to risk (Denburg et al., 2006). The form of the probability weighting

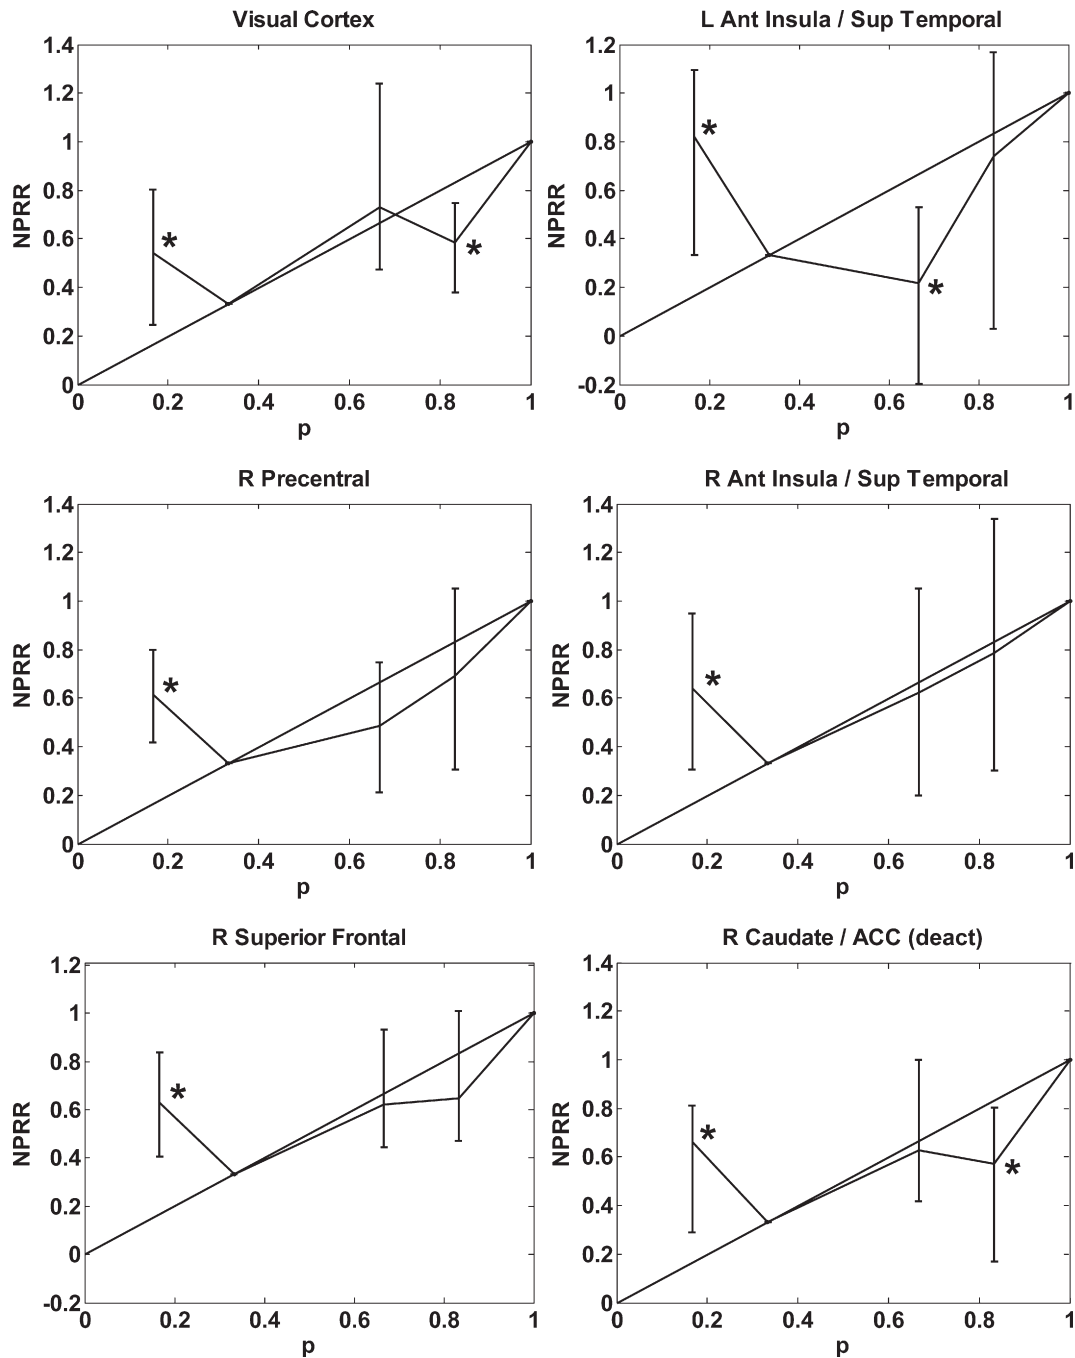


Fig. 3. Neurobiological probability response ratios (NPRR) for representative brain regions in Tables 1 and 2. The mean ratio for each subject was computed at each probability from the three highest voltages, and the median across subjects of the mean value at each probability is shown with error bars indicating the 95% confidence interval for the median ($N=28$). The NPRR curves are plotted with reference to the diagonal, which implies linear probability weighting. Points that are significantly off the diagonal ($P<0.05$) are denoted with “*” and indicate nonlinear probability weighting. Notably, regions associated with visual perception (visual cortex, precuneus) and anterior insula/superior temporal regions displayed significant departures from linear probability weighting. The right caudate/subgenual cingulate showed deactivation to prospects of increasing magnitude but still showed significantly nonlinear probability weighting (*bottom right*).

function derived from the SCR data, however, was predominately linear, suggesting that the source of nonlinear transformation of probabilities may not be in the arousal system as suggested by the somatic marker hypothesis (Damasio, 1999).

Our fMRI data, however, are consistent with a biological bias toward nonlinear weighting of probabilities. In the passive condition, participants were simply presented with a visual cue that

conveyed magnitude and probability information about an impending electrical shock. Because no decision was called for, the brain responses represented an intrinsic reaction to the information presented. Of the regions that responded to the cue, there was a range of relationships to probability, but in almost all cases the biological probability weighting function displayed an “inverted S” shape, which is characteristic of the behavioral probability




Fig. 4. Probability response ratio (NPRR), for integrated skin conductance responses (SCR). Using the form in Eq. (2), the SCR was not significantly different from linear probability weighting across the three highest voltage prospects (median \pm 95% confidence interval shown, $N=26$).

weighting function hypothesized in prospect theory. Regions that have been previously associated with probabilistic decision making, notably in the parietal cortex (Huettel et al., 2006; Platt and Glimcher, 1999; Shadlen and Newsome, 2001), showed significant forms of nonlinearity (Tables 1 and 2). Regions of the parietal cortex, specifically LIP in monkeys, have been proposed to contain a map of the expected utility of all possible actions (Glimcher et al., 2005). Although we did not find a map organized exactly along the dimension of expected utility, we did find maps for both probability and magnitude near the temporo-parietal junction that lay adjacent to each other, if not quite overlapping. The temporo-parietal junction, including both the inferior parietal lobule and superior temporal gyrus has been implicated in at least one previous studies of risky financial decision making (Paulus and Frank, 2006). This region has also been suggested to play a critical role in the judgment of true and false beliefs originating from other people (Grezes et al., 2004; Saxe and Kanwisher, 2003; Sommer et al., in press) as well as attention shifting (Shulman et al., 2007). Its role in our experiment may operate similarly, if more generally outside the specific circumstance of decision making: judging the likelihood of receiving and avoiding an aversive outcome (Mitchell, 2007).

Another region that showed nonlinear weighting of probabilities was a large bilateral cluster encompassing the anterior insula and the superior temporal gyrus. The anterior insula, in particular, has been previously associated with anticipatory responses to painful stimuli (Craig, 2003; Koyama et al., 2005; Ploghaus et al., 1999; Tracey, 2005). Although this region was identified by its monotonically increasing response to expected magnitude, it also showed a significantly nonlinear NPRR (Fig. 3), suggesting that the distortion of probabilities becomes intertwined with representations of the likelihood of subjective states. Notably, we did not find significant activations of the amygdala, which has previously been associated with learning associations between cues and shocks, but this could be due to the extended nature of the cues we used, or to the tendency of the amygdala to be activated transiently during cue acquisition (LaBar et al., 1998).

For the most part, different brain regions were responsive to potential outcome magnitude (i.e. voltage) and probability. Many economic theories of decision-making, such as expected utility or prospect theory, implicitly assume that the individual processes

probability and magnitude separately at some level and that the product of functions of these two terms governs the actual decision. Although anatomical dissociation of these functions, as we found, lends support to the idea that people process these two dimensions separately, there do appear to be subtle interdependencies, which may be dependent on the specifics of the task design. The close interposition of probability and magnitude maps at the temporo-parietal junction suggests a topographic organization of the salient dimensions of the task. However, the significant interaction of these two dimensions was observed in only one region, the ACC. It is notable that the ACC has been previously implicated in modulating decision weights for risky financial decisions (Paulus and Frank, 2006). However, unlike this previous study, we found no significant correlation between individuals' levels of ACC activation and the curvature of their probability weighting implied by their lottery choice decisions. One possible reason is that the activation we measured represented a passive response in the absence of a decision. The ACC is a prime candidate for the integration of magnitude and probability information – even in the absence of a required response, because of its prominent integrative role between the affective system and the motor response system (Botvinick et al., 2001; Critchley et al., 2001; Miller and Cohen, 2001). Although prospect theory assumes functional separability of probability and utility functions, the multiplication of these two dimensions must occur somewhere. LIP has been a promising candidate in monkeys, the ACC may be an additional target in humans. Moreover, no comparable data exist in monkeys using aversive paradigms. The ACC is an integral component of the cortical pain matrix, and current evidence suggests that the ACC integrates several dimensions of the subjective pain experience,

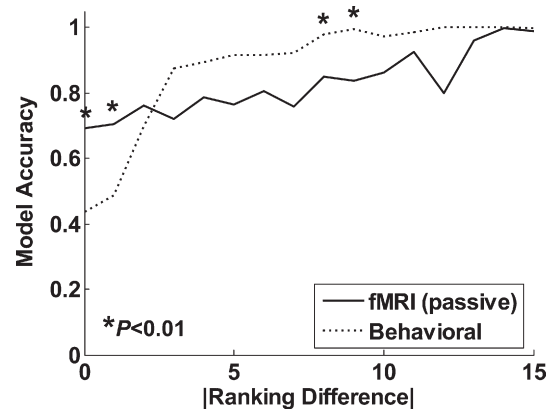


Fig. 5. Accuracy of predicting choice from passive brain activation (solid line) and behavior alone (dashed line). Using a bootstrap logistic regression with inputs of passive activations from regions of interest, we accurately predicted 79% of an individual's choices during the active phase. The predictive accuracy was a function of how similar in rank the choices were. The model accuracy was calculated as the difference in output from the passive brain activation or lottery ranks (from behavioral choices) from the actual choice on trials not used in the regression. Here, we present the results of 10,000 bootstraps across 28 subjects. When the two lotteries were ranked similarly ($|\text{Ranking Difference}| \leq 2$), the predictive accuracy of the passive brain activation was significantly better than choices predicted by the behavioral model alone. When the lotteries were sufficiently different in rank (e.g. $|\text{Ranking Difference}| \geq 3$), the passive brain activation was nearly the same as the behavioral model (except for 2 points as noted). Both models predicted close to 100% of an individual's choices when the rankings were sufficiently different ($|\text{Ranking Difference}| > 10$).

including emotional and attentional components (Craig, 2003; Ploghaus et al., 1999; Vogt, 2005). Interestingly, studies of positive financial incentives have also shown a significant interaction of probability and magnitude in the ACC (Knutson et al., 2005).

In contrast to anticipating painful stimuli, there is the possibility that participants anticipated the pleasure of avoiding a shock. Despite the lack of correlation between the ex-post ratings with voltage or probability in the no-shock outcomes, these ex-post ratings did not fully capture the anticipatory experience. During the cue phase, we did, in fact, observe negative correlations with both expected magnitude and probability in regions classically associated with rewarding stimuli. Notably, the head of the caudate, ventral striatum, and subgenual ACC extending into the orbito-frontal cortex displayed significant negative correlations (Tables 1 and 2). These were all deactivations relative to the implicit baseline, and therefore are somewhat problematic to interpret. An obvious interpretation would be consistent with these regions' roles in reward-anticipation (Knutson et al., 2005). However, a role for salience, which is context-dependent, cannot be ruled out either (Zink et al., 2006). In any case, the NPRR (Eq. (2)) is defined for negative activations, and these regions showed significant departures from linearity in the direction predicted by prospect theory (see Fig. 3).

Taken together, we find that departures from linear probability weighting are seen in several brain systems. These nonlinearities appear at the perceptual (e.g. visual and parietal cortex), visceral (e.g. insula), reward anticipation (e.g. caudate), and motor preparation (premotor and frontal cortex) stages. To the extent that these regions represent serial processing, it would appear that the nonlinearities may arise at the very earliest stage: the perceptual one. In part, this may result from the way in which probabilistic information is presented visually, but this remains a question for future research.

The main result of this paper, however, is the consistent form of nonlinearity observed in neurobiological probability response ratios that parallels the type of nonlinearity observed behaviorally. When observing choice behavior alone, the source of this nonlinearity is unknowable. Individuals, for example, might distort probabilities only when they make a decision. Our data, however, suggests that these distortions occur even in the absence of choice and thus are a property of a more fundamental process of how the brain transforms representations of probabilities into biological responses. Moreover, we conjecture that the nonlinearities at the biological level and in choice behavior may, in fact, be causally connected. This connection is not necessarily a one-to-one mapping from choices to specific brain regions. Recent approaches of utilizing fMRI data to predict an individual's choices have met with varying success rates. In general, approaches that utilize distributed patterns of activation do better than approaches relying on individual brain regions (Hampton and O'Doherty, 2007; Knutson et al., 2007; LaConte et al., 2007). The results of the regression analysis of our choice data indicate that there is a strong relationship between passively evoked activations and subsequent decisions about the same stimuli. However, we observed a high degree of heterogeneity among individuals in their estimated coefficients. Although the same ROIs were used for each individual, the way in which each person represented the probability and voltage information throughout this network was different. Thus, while we observed consistent patterns of probability weighting in the aggregate, individual differences precluded the determination of one canonical set of regression coefficients for

these regions. Even so, the predictive power of the brain activations in individual participants was remarkable and approached 100% for prospects that generated sufficiently different activations.

Even more surprising was the fact that the passive brain activations were significantly better than the choice behavior on its own at predicting choice when the rankings of the available alternatives were similar. This result has two important implications. First, the high degree of predictive accuracy suggests that people possess intrinsic preferences that are stable between passive experience and active decision making (von Neumann and Morgenstern, 1944). Second, the poor predictive power of lottery ranks compared to brain activation when the ranks are similar suggests that the reliance on choice alone to characterize an individual's preferences may give the appearance that choice is probabilistic or that preferences are unstable (Harless and Camerer, 1994; Hey and Orme, 1994; Machina, 1985). The fact that passive brain activations can predict behavior significantly better than the ranking suggests that although choices may be probabilistic or inconsistent, preferences may not be. Choice may be subject to a stochastic process, such as an error in evaluating the payoffs of the available alternatives, which is not present during passive experience (Luce, 1959). Although recent studies have used brain activation to predict decisions (Hampton and O'Doherty, 2007; Knutson et al., 2007), these studies have relied on activation during the choice itself, and may therefore be subject to a similar stochastic component. This would explain why our results have a higher predictive power using passive activations than other studies have found using activation during decision making.

Acknowledgments

We are grateful to Giuseppe Pagnoni and Whitney Herron for assistance and input throughout this experiment. Supported by grants from the National Institute on Drug Abuse (DA016434 and DA20116).

References

- Abdellaoui, M., 2000. Parameter-free elicitation of utility and probability weighting functions. *Manag. Sci.* 46, 1497–1512.
- Allais, M., 1953. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica* 21, 503–546.
- Berns, G.S., Chappelow, J.C., Cekic, M., Zink, C.F., Pagnoni, G., Martin-Skurski, M.E., 2006. Neurobiological substrates of dread. *Science* 312, 754–758.
- Berns, G.S., Capra, C.M., Moore, S., Noussair, C., 2007. A shocking new experiment: new evidence on probability weighting and common ratio violations. *Judgm. Decis. Making* 2, 234–242.
- Botvinick, M.M., Braver, T.S., Barch, D.M., Cohen, J.D., Carter, C.S., 2001. Conflict monitoring and cognitive control. *Psychol. Rev.* 108, 624–652.
- Camerer, C., Loewenstein, G., Prelec, D., 2005. Neuroeconomics: how neuroscience can inform economics. *J. Econ. Lit.* 9–64.
- Colley, W.N., 2002. Colley's bias free college football ranking method: the Colley matrix explained. <http://www.colleyrankings.com/matrate.pdf>.
- Craig, A.D., 2003. Pain mechanisms: labeled lines versus convergence in central processing. *Annu. Rev. Neurosci.* 26, 1–30.
- Critchley, H.D., Mathias, C.J., Dolan, R.J., 2001. Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron* 29, 537–545.
- Damasio, A.R., 1999. *The Feeling of What Happens. Body and Emotion in the Making of Consciousness.* Harcourt Brace and Company, New York.

- Davison, A.C., Hinkley, D.V., 1997. *Bootstrap Methods and Their Application*. Cambridge University Press, Cambridge.
- Denburg, N.L., Recknor, E.C., Bechara, A., Tranel, D., 2006. Psychophysiological anticipation of positive outcomes promotes advantageous decision-making in normal older persons. *Int. J. Psychophysiol.* 61, 19–25.
- Fehr-Duda, H., de Gennaro, M., Schubert, R., 2006. Gender, financial risk, and probability weights. *Theor. Decis.* 60, 283–313.
- Glimcher, P.W., 2002. Decisions, decisions, decisions: choosing a biological science of choice. *Neuron* 36, 323–332.
- Glimcher, P.W., Dorris, M.C., Bayer, H.M., 2005. Physiological utility theory and the neuroeconomics of choice. *Games Econ. Behav.* 52, 213–256.
- Grezes, J., Frith, C.D., Passingham, R.E., 2004. Inferring false beliefs from the actions of oneself and others: an fMRI study. *NeuroImage* 21, 744–750.
- Hampton, A.N., O'Doherty, J.P., 2007. Decoding the neural substrates of reward-related decision-making with functional MRI. *Proc. Natl. Acad. Sci. U. S. A.* 104, 1377–1382.
- Harless, D.W., Camerer, C.F., 1994. The predictive utility of generalized expected utility theories. *Econometrica* 62, 1251–1289.
- Hey, J.D., Orme, C., 1994. Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62, 1291–1326.
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., Camerer, C.F., 2005. Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310, 1680–1683.
- Huettel, S.A., Stowe, C.J., Gordon, E.M., Warner, B.T., Platt, M.L., 2006. Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49, 765–775.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263–291.
- Knutson, B., Taylor, J., Kaufman, M., Peterson, R., Glover, G., 2005. Distributed neural representation of expected value. *J. Neurosci.* 25, 4806–4812.
- Knutson, B., Rick, S., Wimmer, G.E., Prelec, D., Loewenstein, G., 2007. Neural predictors of purchases. *Neuron* 53, 147–156.
- Koyama, T., McHaffie, J.G., Laurienti, P.J., Coghill, R.C., 2005. The subjective experience of pain: where expectations become reality. *Proc. Natl. Acad. Sci. U. S. A.* 102, 12950–12955.
- LaBar, K.S., Gatenby, J.C., Gore, J.C., LeDoux, J.E., Phelps, E.A., 1998. Human amygdala activation during conditioned fear acquisition and extinction: a mixed-trial fMRI study. *Neuron* 20, 937–945.
- LaConte, S.M., Peltier, S.J., Hu, X.P., 2007. Real-time fMRI using brain-state classification. *Hum. Brain Mapp.* 28, 1033–1044.
- Luce, R.D., 1959. *Individual Choice Behavior*. John Wiley and Sons, New York.
- Machina, M.J., 1985. Stochastic choice functions generated from deterministic preferences over lotteries. *Econ. J.* 95, 575–594.
- Miller, E.K., Cohen, J.D., 2001. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
- Mitchell, J.P., 2007. Activity in the right temporo-parietal junction is not selective for theory-of-mind. *Cereb. Cortex* (Electronic publication ahead of print).
- Noussair, C., Robin, S., Ruffieux, B., 2004. A comparison of hedonic rating and demand revealing auctions. *Food Qual. Prefer.* 15, 393–402.
- Paulus, M.P., Frank, L.R., 2006. Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *NeuroImage* 30, 668–677.
- Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Ploghaus, A., Tracey, I., Gati, J.S., Clare, S., Menon, R.S., Matthews, P.M., Rawlins, J.N.P., 1999. Dissociating pain from its anticipation in the human brain. *Science* 284, 1979–1981.
- Preusschoff, K., Bossaerts, P., Quartz, S.R., 2006. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51, 381–390.
- Savage, L.J., 1954. *The Foundations of Statistics*. Wiley, New York.
- Saxe, R., Kanwisher, N., 2003. People thinking about people: the role of the temporo-parietal junction in “theory of mind”. *NeuroImage* 19, 1835–1842.
- Shadlen, M.N., Newsome, W.T., 2001. Neural basis of perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.
- Shulman, G.L., Astafiev, S.V., McAvoy, M.P., d'Avossa, G., Corbetta, M., 2007. Right TPJ deactivation during visual search: functional significance and support for a filter hypothesis. *Cereb. Cortex* 17, 2625–2633.
- Sommer, M., Dohnel, K., Sodian, B., Meinhardt, J., Thoermer, C., Hajak, G., in press. Neural correlates of true and false belief reasoning. *NeuroImage* 35, 1378–1384.
- Starmer, C., 2000. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *J. Econ. Lit.* 332–382.
- Tracey, I., 2005. Nociceptive processing in the human brain. *Curr. Opin. Neurobiol.* 15, 478–487.
- Tversky, A., Kahneman, D., 1992. Advances in prospect theory. Cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297–323.
- Vogt, B.A., 2005. Pain and emotion interactions in subregions of the cingulate gyrus. *Nat. Rev., Neurosci.* 6, 533–544.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton.
- Woodruff, R.S., 1952. Confidence intervals for medians and other position measures. *J. Am. Stat. Assoc.* 47, 635–646.
- Wu, G., Gonzalez, R., 1996. Curvature of the probability weighting function. *Manag. Sci.* 42, 1676–1690.
- Zink, C.F., Pagnoni, G., Chappelow, J., Martin-Skurski, M., Berns, G.S., 2006. Human striatal activation reflects degree of stimulus saliency. *NeuroImage* 29, 977–983.