

Neural Response to Reward Anticipation under Risk Is Nonlinear in Probabilities

Ming Hsu,^{1,2} Ian Krajbich,³ Chen Zhao,⁴ and Colin F. Camerer³

¹Beckman Institute for Advanced Science and Technology and ²Department of Economics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, ³Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, California 91125, and ⁴Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

A widely observed phenomenon in decision making under risk is the apparent overweighting of unlikely events and the underweighting of nearly certain events. This violates standard assumptions in expected utility theory, which requires that expected utility be linear (objective) in probabilities. Models such as prospect theory have relaxed this assumption and introduced the notion of a “probability weighting function,” which captures the key properties found in experimental data. This study reports functional magnetic resonance imaging (fMRI) data that neural response to expected reward is nonlinear in probabilities. Specifically, we found that activity in the striatum during valuation of monetary gambles are nonlinear in probabilities in the pattern predicted by prospect theory, suggesting that probability distortion is reflected at the level of the reward encoding process. The degree of nonlinearity reflected in individual subjects’ decisions is also correlated with striatal activity across subjects. Our results shed light on the neural mechanisms of reward processing, and have implications for future neuroscientific studies of decision making involving extreme tails of the distribution, where probability weighting provides an explanation for commonly observed behavioral anomalies.

Key words: probability weighting; prospect theory; expected utility theory; fMRI; decision making; reward

Introduction

A central question in the social and biological sciences is how the likelihood and value of outcomes are combined to make risky choices. For humans, these choices range widely, from voting, gambling, and buying stocks and insurance to searching for jobs and mates. The standard approach, expected utility theory (EU), assumes that outcomes x are valued nonlinearly by a utility function $u(x)$, but are weighted by their objective probabilities (Savage, 1954).

There is substantial evidence in economics and decision-making research, however, that the hypothesis of expected utility being linear in probabilities is systematically wrong, in a reliable direction. This stylized fact was originally established by Maurice Allais via the “common ratio effect” (Allais, 1953). Consider a decision maker who prefers a sure ($p = 1$) gain of \$100,000 over a coin toss ($p = 0.5$) for \$300,000 but who also rejects a 0.02 chance of \$100,000 for a 0.01 chance of \$300,000. If these choices are consistent with choosing the gamble with the highest expected utility, the first choice implies $u(\$100,000) > 0.5u(\$300,000)$, while the second choice implies $0.02u(\$100,000) < 0.01u(\$300,000)$. The two inequalities are clearly at odds, since both sides of the second are the same as those in the first multiplied by 0.02 (Prelec, 1998).

The common ratio pattern can be reconciled by the plausible assumption that people apply nonlinear “decision weights” $\pi(p)$ to objective probabilities p , so that the ratio $\pi(0.02)/\pi(0.01)$ is much smaller than $\pi(1)/\pi(0.5)$ (cf. Rubinstein, 1988). An inverse S-shaped nonlinear function was first suggested experimentally (Preston and Baratta, 1948), is a central feature of prospect theory (Kahneman and Tversky, 1979), and has been replicated in subsequent experimental and field studies including ours (see Fig. 1A,B and supplemental material, available at www.jneurosci.org). Small probabilities are typically overweighted while high probabilities are underweighted, with a crossover point, at which a probability is subjectively weighted by its objective value of $p^* = \pi(p^*)$ around 1/3 (Kahneman and Tversky, 1979; Kagel et al., 1995; Prelec, 1998; Starmer, 2000).

Despite its clear relevance in a wide variety of settings, few studies have directly studied nonlinear weighting of probabilities, especially compared with the number of studies on risk and reward as a whole. There is much evidence that a number of brain regions are sensitive to expected reward (or “utility”). Arguably the most well established are dopaminergic regions such as the striatum and midbrain structures (Knutson et al., 2001; Abler et al., 2006; Tobler et al., 2007). However, most studies do not sample sufficiently near the probability endpoints to detect the theoretical nonlinearity. The two existing studies on probability weighting also have not shown neural responses to probabilities resembling the smoothly increasing function which typically fit behavior well. Paulus and Frank (2006) focused on between-subjects measures and found that activity in anterior cingulate correlated with degree of nonlinearity across subjects. Berns et al. (2008) used probabilities of shock and found a number of regions

Received Nov. 3, 2008; revised Dec. 15, 2008; accepted Jan. 5, 2009.

This work was supported by a Gordon and Betty Moore Foundation grant and a Human Frontier Science Program grant to C.F.C. We thank P. Bossaerts, K. Friston, A. Healy, and R. Poldrack for comments.

Correspondence should be addressed to Colin F. Camerer, Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125. E-mail: camerer@hss.caltech.edu.

DOI:10.1523/JNEUROSCI.5296-08.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/292231-07\$15.00/0

exhibiting flat responses to probability, but were not able to statistically reject the null linear hypothesis. In this study, we used a parametric design that varied the probabilities and outcome of two gambles in a binary choice task. Our statistical analysis focused on separating the weighted form of expected utility into two components: a portion linear in probabilities, and one that is nonlinear in probabilities. The null hypothesis of EU predicts only significant responses to the linear portion. Under the alternative hypothesis of nonlinear weighting, we expect to find regions correlated with both the linear and nonlinear portions in a manner predicted by existing models of probability weighting.

Materials and Methods

The experimental sample was 21 subjects (11 female). Their mean (SD) age was 29.6 (7.5). Informed consent was given through a form approved by the Internal Review Board at Caltech. Subjects were recruited from an online bulletin board (see supplemental material, available at www.jneurosci.org).

Experimental paradigm

The experiment consisted of 120 self-paced trials (Fig. 2). In each trial, subjects chose between two simple gambles, (p_1, x_1) and (p_2, x_2) . A gamble pays off \$ x with probability p and \$0 otherwise. Subjects were first presented a screen containing only (p_1, x_1) . This serves to isolate the brain response evaluating (p_1, x_1) without confounding evaluation of the second gamble (p_2, x_2) or the process of choice. The values of p_1 are {0.01, 0.1, 0.3, 0.5, 0.8, 0.95}, the values of \$ x are {10, 20, 50, 100}, and they are combined factorially to form 24 pairs, which are each presented five times.

On 12 randomly chosen trials (10% of the trials), subjects were asked after the first-gamble presentation to indicate whether $p_1 > 0.4$. This was done to ensure that subjects were paying attention to the first-gamble stimulus. The final screen presented both the first gamble and a second gamble, and subjects chose one of the gambles. The probability and reward levels of the second gamble, (p_2, x_2) , were varied across trials, and were chosen so that its weighted expected utility (see model below) would be close to that of the first gamble, which facilitated powerful behavioral estimation of the relevant parameters of the experiment. At the end of the experiment, one choice round was randomly selected and the gamble chosen in that round was resolved to determine the subject's payment. Average earnings for subjects were $\$21.48 \pm \5.49 plus the participation fee of \$20 (for more detail, see supplemental Methods, available at www.jneurosci.org as supplemental material).

Behavioral analysis

A stochastic choice model was used to infer the probability weighting function from behavior and correlate its parametric expression with




Figure 1. Nonlinear weighting of probability inferred from choices. **A**, Fits of the weighting function $\pi(p)$ from many previous behavioral studies (see supplemental material, available at www.jneurosci.org as supplemental material). **B**, Fits from individual subjects in our experiment using the one-parameter Prelec weighting function (with $\pi(p) = p$ at $p = 1/e$). **C**, Fits from various weighting functions (supplemental Table S2, available at www.jneurosci.org as supplemental material) using group-level parameters from our experiment (blue: Prelec 1-parameter; red: Prelec 2-parameter; yellow: Kahneman and Tversky; cyan: Lattimore; green: Wu and Gonzalez).

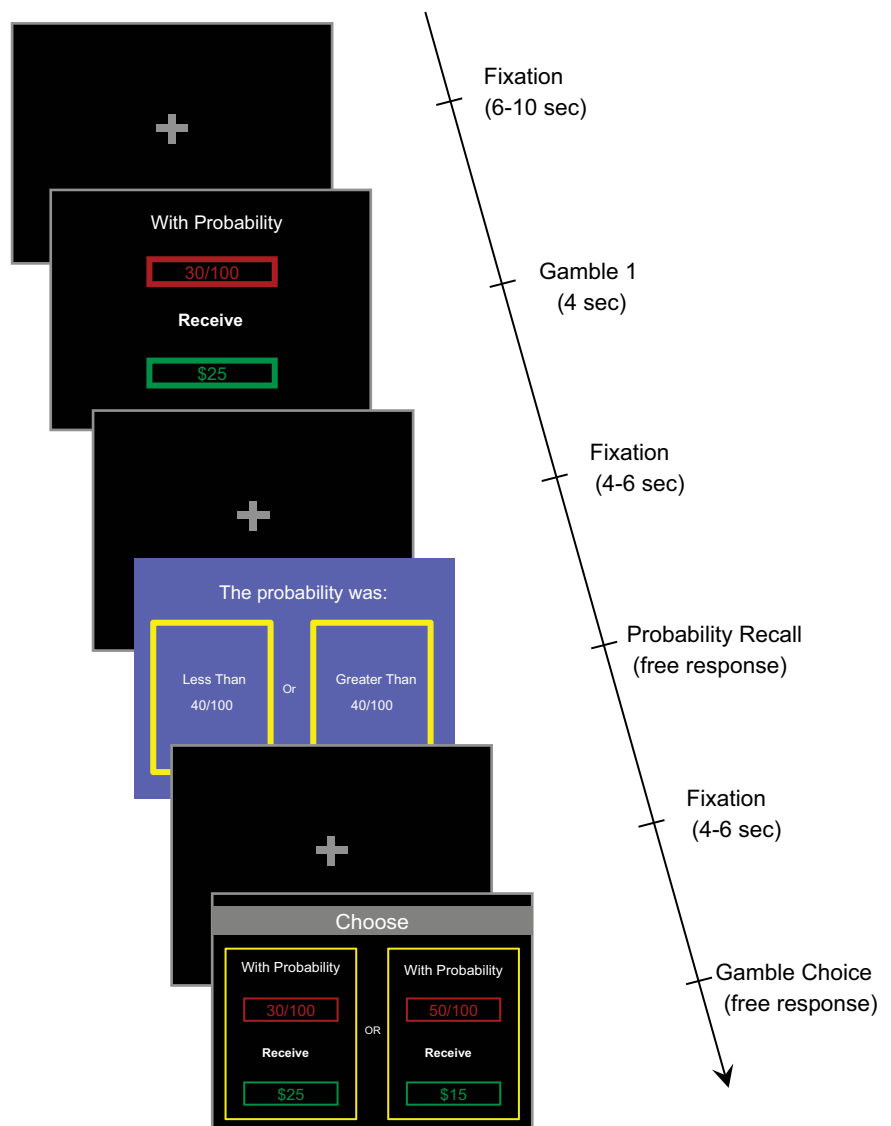


Figure 2. The experimental sequence of events. (1) A single gamble, consisting of the probability p_1 of receiving some dollar amount \$ x_1 (or 0 otherwise). (2) In 12 of the 120 trials, subjects are then asked to indicate whether the probability in the previous screen was greater or less than 40/100 (to engage attention to screen 1). (3) Subjects see a choice screen showing the gamble shown in 1 and a new gamble.

Table 1. Pooled estimates (SEs) of various weighting functions defined in supplemental Table S1 (available at www.jneurosci.org as supplemental material)

	LL	ρ	θ_1	θ_2	λ
Prelec 1 ($\theta_1 = \alpha$)	−1493.4	0.499 (0.0160)	0.724 (0.0235)		1.464 (0.0911)
Prelec 2 ($\theta_1 = \alpha, \theta_2 = \beta$)	−1491.1	0.441 (0.0307)	0.708 (0.0222)	0.859 (0.0619)	1.972 (0.2908)
Kahneman and Tversky ($\theta_1 = \gamma$)	−1505.1	0.494 (0.0161)	0.774 (0.0215)		1.386 (0.0819)
Lattimore ($\theta_1 = \gamma, \theta_2 = \delta$)	−1502.3	0.531 (0.0264)	0.786 (0.0368)	0.795 (0.0592)	1.166 (0.1326)
Wu and Gonzalez ($\theta_1 = \gamma, \theta_2 = \alpha$)	−1504.0	0.525 (0.0274)	0.850 (0.0595)	2.023 (0.7828)	1.191 (0.1446)

brain activity (see supplemental Methods, available at www.jneurosci.org as supplemental material). As in many studies on this topic, precision is gained by assuming specific functional forms for the utility and probability weighting functions. The utility function is assumed to be a power function $u(x; \rho) = x^\rho$, and the probability weighting function is the one-parameter Prelec weighting function $\pi(p; \alpha) = 1/\exp\{\ln(1/p)\}^\alpha$ (Prelec, 1998), derived axiomatically from psychophysical principles [it is implied if $\ln(1/\pi(p))$ is a power function of $\ln(1/p)$]. Decision weights and utilities are assumed to be combined linearly, $U(x, p) = \pi(p, \alpha)u(x, \rho)$ (where the zero payoff has zero utility). We chose the one-parameter Prelec function because it fits as well as or better than other one- and two-parameter functions that were estimated from subjects' choices (Table 1), and having one parameter simplifies the cross-subject analysis below.

fMRI acquisition

Scans were acquired using the 3 Tesla Siemens Trio scanner at Caltech's Broad Imaging Center. Anatomical images (high-resolution, T1-weighted) were acquired first. Functional (T2-weighted) images were then acquired using the following parameters: TR = 2000 ms, TE = 40 ms, slice thickness = 4 mm, 32 slices. Horizontal slices were acquired $\sim 30^\circ$ clockwise of the anterior–posterior commissure (AC–PC) axis to minimize signal dropout of the orbitofrontal cortex. The total time duration of the experiment varied because each round was self-paced.

fMRI analysis

Imaging data were preprocessed using SPM2, including, in order, slice time correction, motion correction, coregistration, normalization to the MNI template, and smoothing of the functional data with an 8 mm kernel (Friston et al., 1995). Random effects analyses were done in SPM2 (Friston et al., 1995) by specifying a separate general linear model for each subject and pooling at the second level. First all images were high-pass filtered in the temporal domain (filter width 128 s) and autocorrelation of the hemodynamic responses was modeled as an AR(1) process. In the GLM model all visual stimuli and motor responses were entered as separate regressors that were constructed by convolving a hemodynamic response function (hrf) with a comb of Dirac functions at the onset of each visual stimulus or motor response. Parametric modulators were added to the main regressors as interaction terms.

Parametric model. An event-related analysis focused on brain activity during presentation of the first gamble. Because no information is present regarding the second gamble, we assume brain regions correlated with decision variables reflect reward anticipation with respect to first gamble, rather than choice. We further make the assumption that neural activity is approximately a linear function of the behaviorally derived utility function (that is, we search for brain activity which resembles closely the functions in Fig. 1A,B). A GLM is used that separates the weighting function into two components: (1) component that is linear in p and (2) the component that is the nonlinear deviation term $\Delta(p, \alpha_i) = \pi(p, \alpha_i) - p$ (Fig. 3A). Specifically, we are looking for a prospect-theoretic expected value function that is nonlinear in p ; that is, $\pi(p, \alpha)u(x) = p \cdot u(x) + \Delta(p, \alpha) \cdot u(x)$. We assume the function $u(x)$ is a power function x^ρ , where the value of ρ is taken from the individual behavioral estimate, and $\Delta(p, \alpha) = \pi(p, \alpha) - p$, where the mean group $\alpha = 0.771$ is used.

The BOLD signal during the first gamble presentation is regressed against $p \cdot u(x)$ and $\Delta(p, \alpha) \cdot u(x)$. If the expected utility (EU) null hypothesis is an accurate approximation of valuation of risky choices, there should be no reward-related brain regions that respond to the deviation term $\Delta(p, \alpha) \cdot u(x)$. If the nonlinear weighting hypothesis is an accurate approximation, there should be reward-related brain regions

that respond equally strongly to the linear component $p \cdot u(x)$ and to the nonlinear component $\Delta(p, \alpha) \cdot u(x)$.

Nonparametric model. To see how closely activity in brain regions correlated with the weighting function corresponds to the behaviorally derived stylized empirical weighting function, a nonparametric method was used. Each level of probability p was given a separate dummy variable $I(p)$. The relation $y = \alpha + \beta_i \cdot I_i(p) \cdot u(x) + \varepsilon$ was estimated, where y is the BOLD response upon presentation of the first gamble, I is an indicator function for the particular level of p , and $u(x)$ is a power function x^ρ used above. Each β_i for the six levels of probability was then rescaled by dividing the β_i values for each probability level by the estimated slope for the response to the linear probability term in the parametric regression of activity against the linear and deviation terms. This is to take into account the dimensionless nature of BOLD responses as well as to add a robustness check of the concordance of the relationship between the parametric and nonparametric model.

Between-subject correlation. Next we test whether cross-subject variation in the inflection of nonlinear weighting inferred from choices is consistent with cross-subject differences in neural activity. Intuitively, more highly nonlinear functions will be approximated by a combination of the linear term p and the nonlinear term $\Delta(p, \alpha_i) = \pi(p, \alpha_i) - p$ (shown in Fig. 3A, right) that puts more weight on the nonlinear term. Less nonlinear functions will put less weight on the nonlinear term. A linear-weighting subject, for example, will put no weight on the nonlinear deviation $\Delta(p, \alpha_i) = \pi(p, \alpha_i) - p$. This analysis exploits the helpful fact that in the one-parameter Prelec form, all functions $\pi(p, \alpha_i)$ pass through a common value $\pi(1/e) = 0$. Weighting functions with $\alpha > 0.77$ (less inflected, more linear) will therefore be well approximated by p plus a dampened form of the deviation curve in Figure 3A (right). Weighting functions with $\alpha < 0.77$ (more inflected) will be approximated by p plus an amplified form of the decision curve in Figure 3A.

Denote the true weighting function for subject i by $\pi(p, \alpha_i)$, and the deviations from linear weighting by $\Delta(p, \alpha_i) = \pi(p, \alpha_i) - p$. A brain region that represents $\pi(p, \alpha_i)$ will be significantly correlated with both $\Delta(p, \alpha_i)$ and p . Using the mean inflection parameter $\bar{\alpha} = 0.77$, the theoretical relation is $\pi(p, \alpha) = a + b_1 \cdot p + b_2 \cdot \Delta(p, \bar{\alpha}) + e$ [that is, every weight α gives a weighting function that, when its values are regressed against p and $\Delta(p, \bar{\alpha})$, gives weights b_1 and b_2]. Fig. 4A shows the value of b_2 that is theoretically estimated for the various values of α inferred from subjects' choices. The value of b_2 is the inflection sensitivity (compared with the benchmark group mean $\alpha = 0.77$) for particular values of subject-specific α . Note that when $\alpha = \bar{\alpha}$ (a person's α equals the group mean), $b_2 = 1$, and when $\alpha = 1$ (a linear weighting function), $b_2 = 0$ because the nonlinear deviation term receives no weight. The graph shows that in theory, lower individual values of α , corresponding to more inflection of the weighting function, should lead to higher values of the b_2 coefficient. That is, more inflected functions are best approximated by a combination of the linear term p and nonlinear deviation term (calculated using the average $\bar{\alpha}$) with a higher weight on the nonlinear deviation term.

Results

Behavioral results

Table 1 contains the pooled (grouped-subject) estimates of the various weighting function parameters as well as the utility function power parameter ρ and the stochastic choice response sensitivity λ . Five subjects always chose the gamble with the higher probability, such that their choices did not permit identification of the relevant parameters. These subjects were therefore ex-

cluded from this and all subsequent fMRI analyses that depended on these parameter estimates. Figure 1C plots the estimated functional forms for all five functions and shows that they look similar (except for more pronounced overweighting of low probabilities in the Prelec two-parameter form). The two Prelec functions fit slightly better than the others as judged by lower negative log likelihood LL. We focused on the one-parameter version of the Prelec function, as it fits only slightly worse than the two-parameter version, and permits simpler cross-subject comparison (since each subject's curvature is expressed by one parameter rather than two).

Supplemental Table S4 (available at www.jneurosci.org as supplemental material) presents individual parameter estimates for the stochastic choice model. These parameter estimates are similar to those found in the literature, indicating nonlinear weighting of probability (mean $\alpha = 0.77 \pm 0.08$) and concavity of utility for money ($\rho = 0.57 \pm 0.04$) (see supplemental material, available at www.jneurosci.org). These estimates also correspond closely to pooled estimates shown in Table 1.

fMRI results

We first identify regions that are significantly correlated with the linear term of the parametric model (Table 2). Consistent with previous literature, we found a number of regions including striatum (Knutson et al., 2001; Abler et al., 2006; Preuschoff et al., 2006; Tobler et al., 2007), anterior cingulate cortex (ACC) (Knutson et al., 2001), and cerebellum (Knutson et al., 2001), as well as frontal regions such as motor areas and medial prefrontal cortex (Knutson et al., 2001; Yacubian et al., 2006; Tobler et al., 2007). Regions that are significantly correlated with the nonlinear term include the striatum, cingulate gyrus, motor cortex, and cerebellum (Table 3). Significant activation in the striatum and cingulate gyrus are in particular consistent with findings from two previous papers on probability weighting (Paulus and Frank, 2006; Berns et al., 2008). In contrast, there are no regions where activity was negatively correlated with the nonlinear term at traditional significance levels ($p > 0.1$) (for details, see Table S6, available at www.jneurosci.org as supplemental material).

Next we search for regions that are significantly and equally activated by both the linear term and the nonlinear term. These regions include the striatum (Fig. 3B; supplemental Figs. S2, S3, available at www.jneurosci.org as supplemental material), motor cortex, and cerebellum. We focus on the striatum for two reasons. First, it is well established to be involved in reward processing and specifically in reward anticipation (Knutson et al., 2000; Schultz, 2000; O'Doherty et al., 2004). In particular it has been suggested

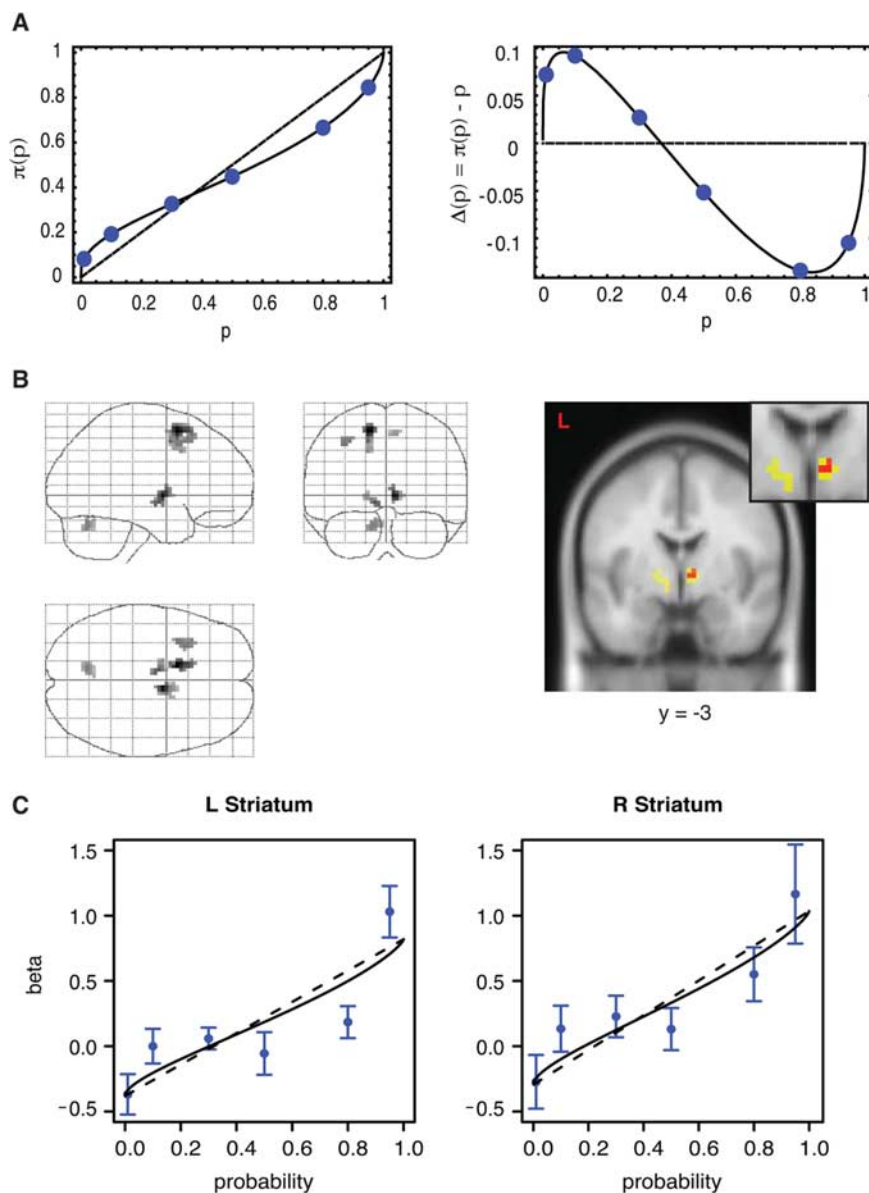


Figure 3. *A*, The analysis decomposes expected reward responses into two terms, the linear component in p (left, dashed line) and the nonlinear $\pi(p) - p$ component (right). *B*, Glass brain and coronal section of activations to both linear p and nonlinear $\pi(p) - p = 0.77$. Red, Regions where both linear and nonlinear terms are activated at $p < 0.001$, excluding regions where linear and nonlinear terms are significantly different at $p < 0.9$; Yellow, Regions where both linear term is activated at $p < 0.001$ and nonlinear term at $p < 0.005$, excluding regions where linear and nonlinear terms are significantly different at $p < 0.5$. For additional coronal sections, see supplemental Figure S2 (available at www.jneurosci.org as supplemental material). *C*, Normalized GLM β coefficients of BOLD signal activation for extracted voxels (blue dots) in the left and right striatum coronal section shown in *B* (yellow) and Prelec function with group behavioral parameter ($\alpha = 0.77$) inferred from choices (solid black line).

that the striatum combines reward magnitude and probability multiplicatively into a general expected reward signal (Tobler et al., 2007). Second, in our data, the striatum is correlated with both the linear term and the nonlinear term (Fig. 3B). We cannot reject the hypothesis of equal response to both terms at highly liberal p values (i.e., we can conclude with high confidence that the activity levels are not different). Furthermore, this result is robust to two variations of the statistical model that we used (see supplemental Figs. S3, S4, available at www.jneurosci.org as supplemental material).

Figure 3C contains results from the nonparametric model, where the neural β normalized by the group mean β for each level of probability used in the experiment are plotted against the true

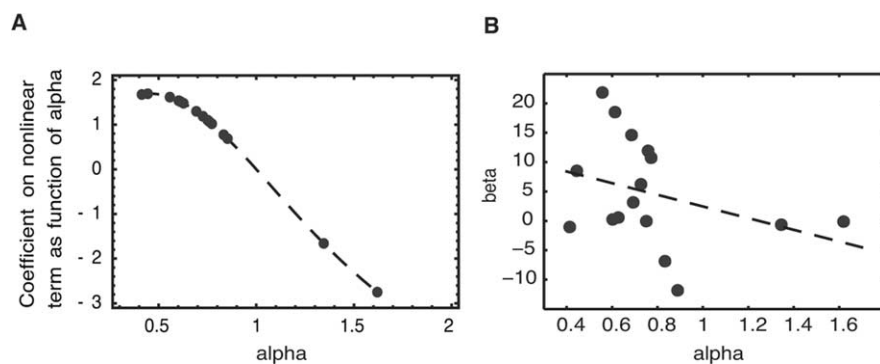


Figure 4. *A*, Theoretical coefficient on nonlinear probability component $\pi(p) - p$ as a function of individual nonlinearity parameter α . *B*, Scatter plot of within-subject empirical response (β coefficient from GLM) in average of left and right striatum and individual-subject nonlinearity parameter α . The negative correlation is consistent with the theoretical relationship shown in *A*.

probabilities, as well as a plot of the behaviorally inferred weighting function using the group $\alpha = 0.77$. To adjust the raw probability weight values (which range from 0 to 1) to scale comparably to the neural β , the behavioral weights were regressed against the neural β for each probability. The behavioral weights were then multiplied by the regression slope, and the regression constant was added, to create the adjusted values plotted in Figure 3C. There is a clear concavity in low-probability activity and convexity in high-probability activity, and relatively equal activity for probabilities from 0.10 to 0.80, which are also the key features of the psychometric curves derived from behavior. The neurometrically derived BOLD signal curve looks like a more inflected expression of the psychometric curve. In particular, the behavioral weighting function departs from the objective probabilities more so than the $I(p)$ neural estimates. Finally, a non-nested model test (Cox test) rejects the linear model in favor of the nonlinear $\pi(p)$ ($p < 0.064$) and does not reject the linear model against the $\pi(p)$ alternative ($p < 0.302$) (see supplemental Methods, available at www.jneurosci.org as supplemental material). That is, the $\pi(p)$ model contains additional explanatory power relative to the linear model, whereas the reverse is not true. This, however, cannot rule out other models that we have not considered, nor potential confounds that may influence the BOLD response to anticipated reward. Indeed, introspection of Figure 3C shows that there remains much variation between the neurometric and psychometric curves that are not explained by our model; the highly imperfect fit could surely be improved by other designs and models. Future studies will therefore be needed to explore this important issue.

Cross-subject correlation

Figure 3B presents the regions of striatum with activity during the first gamble presentation that is significantly correlated with the linear component p and the nonlinear component $\Delta(p, \alpha)$ and in which the empirical coefficient on the nonlinear component (b_2) is negatively correlated with subject α . The values of the nonlinear component weight b_2 for each subject is plotted against that subject's behaviorally estimated α in Figure 4B. Across subjects, those with more inflected decision weights as revealed by the behavioral stochastic choice function (lower α) do have higher values of b_2 estimated from neural activity ($r = -0.35$, bootstrapped 95% confidence interval $(-0.60, -0.05)$). The scaling of the coefficients is arbitrary (since it is derived from BOLD signal) but the sign of the coefficients is consistent with the theoretical relationship in Figure 4A. Subjects with positive b_2 coefficients are generally those with inverse S-shaped weighting functions ($\alpha < 1$), who overestimate small probabilities and underestimate

large ones. The two subjects with negative b_2 coefficients have $\alpha > 1$ and underestimate small probabilities and overestimate large ones [i.e., they act as if they have negative rather than positive weight on the nonlinear deviation term shown in Fig. 3A (right)]. Excluding the two subjects with $\alpha > 1$ from the Figure 4B analysis reduces the correlation coefficient only slightly, to $r = -0.34$ [bootstrap 95% confidence interval $(-0.92, 0.17)$].

Discussion

The hypothesis that organisms weight probabilities objectively (i.e., linearly) when evaluating risks has been a useful benchmark in many areas of social and biological sciences and is a reasonable approximation for many risks (Kagel et al., 1995). However, much of the appeal of the linear-weighting EU model comes from its empirical superiority to the simpler expected value model [where $u(x) = x$] and the intuitive appeal of its logical axioms.

There is much behavioral evidence, however, that linearity appears to break down for very high and low probabilities in a systematic manner (Allais, 1953; Tversky and Kahneman, 1992). Indeed, Oskar Morgenstern, one of the founders of expected utility theory, speculated that linearity is a plausible assumption within the intermediate range of probabilities because “a normal individual would have some intuition of what 50:50 or 25:75 means,” but that linearity at extreme probabilities was unlikely because “probabilities used must be within certain plausible ranges and not go to 0.01 or even less to 0.001” (Morgenstern, 1979, p. 178). Allowing for such nonlinearity elegantly reconciles the common phenomenon of simultaneously purchasing insurance against rare disasters (a risk-averse choice) and buying lottery tickets (a risk-seeking choice) because both are consistent with overweighting low probabilities.

The linearity hypothesis has also been adopted, often implicitly, in most studies of decision making and learning in neuroscience (Yacubian et al., 2006). For example, in standard models of reinforcement learning, reward prediction is assumed to result from an unbiased representation of rewards accumulated from experience (Montague et al., 2004). Therefore, reward predictions in stochastic environments are expected to be accurate in the long run. Most existing studies have not sampled closely enough to probability endpoints to be sensitive to the full pattern of predicted nonlinearity. For example, in the study by Abler et al. (2006), the probabilities sampled were $\{0, 0.25, 0.5, 0.75, 1\}$, for money rewards. Berns et al. (2008) used probabilities $\{1/6, 2/6, 4/6, 5/6, 1\}$ for electric shocks. Given existing evidence that the inflection point is $\sim 1/3$, one would expect to find a closely linear approximation across the range of probabilities in the former studies, rather than the full reverse S-shaped function. Similarly, one would expect a convex function in the study by Berns et al. (2008) with disproportionate brain responses to the lowest probability of 1/6. This is consistent with their results, which found a U-shaped response in the caudate/subgenual anterior cingulate cortex. This region of activation is further anterior compared with our striatal activation, and is negatively correlated with the magnitude of shock. This potentially reflects differences between encoding of rewards in the gain and loss domain.

In addition, in many neuroscientific studies of decision making, probabilities are estimated, either from ratios of bars and pie

charts, or from verbal descriptions of probabilities. This method of estimation is likely to introduce bias into the data, particularly for low probabilities, which are often estimated as higher than they are in reality.

Table 2. Regions significantly correlated with probability (p) term

Voxels (k)	Cluster p_{unc}	t value	Voxel p_{unc}	MNI coordinates			Region
				x	y	z	
159	0.003	5.39	0	−12	15	60	Motor cortex
		5.18	0	−15	9	48	
19	0.236	4.89	0	−12	−3	−6	Striatum
78	0.025	4.76	0	30	−51	−27	Culmen
12	0.345	4.63	0	−9	57	24	Superior frontal gyrus
29	0.147	4.48	0	12	0	−3	Globus pallidus
14	0.307	4.44	0	30	−87	−9	Inferior occipital gyrus
22	0.203	4.4	0	6	24	−12	Brodman 25
26	0.168	4.28	0	−15	−48	24	Posterior cingulate gyrus
		3.78	0.001	−24	−48	27	
41	0.089	4.25	0	−9	−75	−24	Cerebellum
		4.03	0	−9	−66	−27	
17	0.261	4.23	0	−42	9	21	Inferior frontal gyrus
15	0.291	4.14	0	51	18	24	IFG
28	0.154	4.04	0	18	21	18	ACC
		3.89	0.001	27	24	18	
11	0.366	4.03	0	21	−45	15	Brodman 31
		3.81	0.001	18	−48	24	
19	0.236	3.96	0	−18	−6	3	Striatum
		3.95	0	−24	−6	12	
		3.67	0.001	−27	3	9	
13	0.325	3.85	0.001	−33	9	−9	
18	0.248	3.85	0.001	−9	9	−3	Insula/putamen
		3.82	0.001	−21	18	−12	
		3.75	0.001	−18	18	−3	

Activations are thresholded at $p < 0.001$ and cluster size $k > 10$. Coordinates are in MNI space.

Table 3. Regions significantly correlated with nonlinear deviation term $\Delta(p)$

Voxels (k)	Cluster p_{unc}	t value	Voxel p_{unc}	MNI coordinates			Region
				x	y	z	
169	0	5.4	0	33	3	24	Brodman 9/cingulate
		4.65	0	15	24	30	
		4.33	0	27	18	30	
65	0.007	5.06	0	3	−3	0	Striatum
		4.84	0	0	−6	−9	
		4.58	0	0	−15	−3	
22	0.091	4.89	0	18	−21	39	Brodman 31
15	0.156	4.52	0	−18	−72	−33	Cerebellum
21	0.098	4.43	0	−12	12	21	Cingulate
19	0.114	4.27	0	24	−9	18	Striatum
28	0.06	4.14	0	−30	−9	39	Middle frontal gyrus
		3.96	0	−30	6	39	
19	0.114	4.14	0	−15	−3	48	Brodman 24
17	0.133	4.06	0	−24	27	48	Middle frontal gyrus
		3.86	0.001	−18	36	48	Brodman 8

Activations are thresholded at $p < 0.001$ and cluster size $k > 10$. Coordinates are in MNI space.

areas (Huettel et al., 2006; Berns et al., 2008) or from past experience in studies of learning (O'Doherty et al., 2004; Haruno and Kawato, 2006). This methodology differs from those in the behavioral literature on probability weighting, where the probabilities are usually given explicitly in numerical format (and sometimes graphically as well) (Camerer and Ho, 1994; Wu and Gonzalez, 1996; Kilka, 2001). Using devices such as pies or bars to depict probabilities without numerical representations can potentially induce difficulties in interpretation of the source of probability weighting. There is substantial evidence that individuals make systematic errors in proportion estimation, especially for proportions close to zero or one (Varey et al., 1990; Hollands and Dyre, 2000). Therefore it becomes unclear whether the behavioral patterns come from errors in estimating probabilities from graphical displays, or from nonlinearity in the valuation process.

In this study, therefore, we closely followed the representation of decisions under risk used in the behavioral economics literature, while at the same time take into account of the limitations of neuroimaging measures by separating presentation of the gambles in time. This allowed us to dissociate neural contributions of the evaluation of individual gambles from activity due to comparison and choice (Fig. 2). The need for temporal separation of evaluation and choice does not arise in behavioral studies in which the only observable variable is choice. Existing studies do not always distinguish between perception of value and the decision process (Paulus and Frank, 2006). This distinction, however, is of potential importance when using neuroimaging data. The modeling of probability weighting assumes that the value function itself is "distorted" (nonlinear) in probabilities, whereas the decision process is unbiased. The temporal separation therefore allows us to focus on reward perception, rather than simultaneous perception

and choice (Preuschoff et al., 2006). Our fMRI results suggest that striatum activity in evaluation of risks is nonlinear in probabilities in isolation of choice, consistent with standard interpretations of the weighting function. Activity in the striatum is also found by Tom et al. (2007) in response to gamble monetary values, and can be used to link the neural aversion to loss (compared with equal-sized gain) to behavioral loss-aversion.

More generally, this type of study represents an empirical competition between models of risky choices which are rooted in logical axioms (chiefly the independence axiom stated earlier) and models rooted in psychophysics which can be further grounded in evolutionary adaptation. This study completes the exploratory studies to the key elements of prospect theory, the others being loss aversion (Tom et al., 2007) and framing/reference point (De Martino et al., 2006) (for review, see Fox and Poldrack, 2008). Our intuition is that brain activity during valuation of risks is more likely to correspond to the cognitive components of prospect theory than to EU, and it will also be easier to construct an adaptationist account of how evolution would have shaped brains to follow prospect theory rather than EU (Robson, 2002), since prospect theory follows from psychophysics and EU from normative logic. Establishing a neural and evolutionary basis of prospect theory could provide an illustrative example of how the foundation for principles guiding social science might be usefully shifted from relying largely on logic, to respecting biological implementation (which might, of course, include convergence to logical principles as a result of learning or higher-order cognition).

There are a number of more general implications for neuroscientific studies of reward and decision making. The maturation of decision neuroscience and neuroeconomics will likely lead to increasing emphasis on problems involving extreme tails of the distribution (e.g., public fear about rare catastrophic risks, pathological gambling, broad participation in long-tailed lotteries, and preferences for financial assets with positive skewness) (Barberis and Huang, 2008). The numerical ratio of overweighting of low probabilities is dramatic for the functions we estimate: probabilities of 10^{-2} , 10^{-6} , and 10^{-9} are overweighted by factors of 4, 520, and 33,000, respectively.

In addition, some studies suggest that probabilities learned through experience do not exhibit the same type of patterns of behavior as those represented abstractly (Hertwig et al., 2004; Fox and Hadar, 2006). Furthermore, because the learning process is still under investigation, it is unknown how the brain updates probabilities conditional on past events. More studies of the neural basis of response to probabilities represented abstractly, and those learned from experience, are therefore needed to provide a unified framework to understand and reconcile these results.

References

- Abler B, Walter H, Erk S, Kammerer H, Spitzer M (2006) Prediction error as a linear function of reward probability is coded in human nucleus accumbens. *Neuroimage* 31:790–795.
- Allais M (1953) Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* 21:503.
- Barberis N, Huang M (2008) Stocks as lotteries: the implications of probability weighting for security prices. *Am Econ Rev* 98:2066–2100.
- Berns GS, Capra CM, Chappelow J, Moore S, Noussair C (2008) Nonlinear neurobiological probability weighting functions for aversive outcomes. *Neuroimage* 39:2047–2057.
- Camerer CF, Ho TH (1994) Violations of the betweenness axiom and non-linearity in probability. *J Risk Uncertain* 8:167–196.
- De Martino B, Kumaran D, Seymour B, Dolan RJ (2006) Frames, biases, and rational decision-making in the human brain. *Science* 313:684–687.
- Fox CF, Poldrack R (2008) Prospect theory and the brain. In: *Neuroeconomics: decision making and the brain* (Glimcher P, Camerer CF, Fehr E, Poldrack RA, eds), pp 145–174. New York: Elsevier.
- Fox CR, Hadar L (2006) “Decisions from experience” = sampling error + prospect theory: reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment Decision Making* 1:159–161.
- Friston KJ, Holmes AP, Worsley K, Poline JB, Frith CD, Frackowiak RSJ (1995) Statistical parametric maps in functional brain imaging: a general linear approach. *Hum Brain Mapp* 2:189–210.
- Haruno M, Kawato M (2006) Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. *J Neurophysiol* 95:948–959.
- Hertwig R, Barron G, Weber EU, Erev I (2004) Decisions from experience and the effect of rare events in risky choice. *Psychol Sci* 15:534–539.
- Hollands JG, Dyre BP (2000) Bias in proportion judgments: the cyclical power model. *Psychol Rev* 107:500–524.
- Huettel SA, Stowe CJ, Gordon EM, Warner BT, Platt ML (2006) Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49:765–775.
- Kagel JH, Battalio RC, Green L (1995) Economic choice theory: an experimental analysis of animal behavior. New York: Cambridge UP.
- Kahneman D, Tversky A (1979) Prospect theory—analysis of decision under risk. *Econometrica* 47:263–291.
- Kilka M (2001) What determines the shape of the probability weighting function under uncertainty? *Manage Sci* 47:1712.
- Knutson B, Westdorp A, Kaiser E, Hommer D (2000) fMRI visualization of brain activity during a monetary incentive delay task. *Neuroimage* 12:20–27.
- Knutson B, Adams CM, Fong GW, Hommer D (2001) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J Neurosci* 21:RC159(1–5).
- Montague PR, Hyman SE, Cohen JD (2004) Computational roles for dopamine in behavioural control. *Nature* 431:760–767.
- Morgenstern O (1979) Some reflections on utility theory. In: *EU hypotheses and the Allais paradox* (Allais O, Hagen O, eds), pp 175–183. Dordrecht: Reidel.
- O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304:452–454.
- Paulus MP, Frank LR (2006) Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *Neuroimage* 30:668–677.
- Prelec D (1998) The probability weighting function. *Econometrica* 66:497–527.
- Preston MG, Baratta P (1948) An experimental study of the auction-value of an uncertain outcome. *Am J Psychol* 61:183–193.
- Preuschoff K, Bossaerts P, Quartz SR (2006) Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51:381–390.
- Robson AJ (2002) Evolution and human nature. *J Econ Perspect* 16:89–106.
- Rubinstein A (1988) Similarity and decision-making under risk (is there a utility theory resolution to the Allais paradox?). *J Econ Theory* 46:145–153.
- Savage LJ (1954) *The foundations of statistics*. New York: Wiley.
- Schultz W (2000) Multiple reward signals in the brain. *Nat Rev Neurosci* 1:199–207.
- Starmer C (2000) Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *J Econ Lit* 38:332–382.
- Tobler PN, O'Doherty JP, Dolan RJ, Schultz W (2007) Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *J Neurophysiol* 97:1621–1632.
- Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. *Science* 315:515–518.
- Tversky A, Kahneman D (1992) Advances in prospect-theory—cumulative representation of uncertainty. *J Risk Uncertain* 5:297–323.
- Varey CA, Mellers BA, Birnbaum MH (1990) Judgments of proportions. *J Exp Psychol Hum Percept Perform* 16:613–625.
- Wu G, Gonzalez R (1996) Curvature of the probability weighting function. *Manage Sci* 42:1676–1690.
- Yacubian J, Gläscher J, Schroeder K, Sommer T, Braus DF, Büchel C (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J Neurosci* 26:9530–9537.